

# Equilibrium Transition from Loss-Leader Competition: How Advertising Restrictions Facilitate Price Coordination in Chilean Pharmaceutical Retail

Yu (Jasmine) Hao\*

Faculty of Business and Economics  
University of Hong Kong  
haoyu@hku.hk

July 1, 2026

## Abstract

Between December 2007 and April 2008 Chile's three retail pharmacy chains coordinated price increases on 220 medicines, weeks after advertising restrictions ended a comparative-price war that had driven prices below cost. I develop a dynamic model of the transition and estimate its demand from the data. Two forces sustained the war: comparative-price ads broadcast who was cheapest, so undercutting paid; and a coordinated increase holds only if rivals expect it to be matched, but through the war none did. The ban moves both. By collapsing price sensitivity ( $\hat{\alpha}$  from 0.103 to 0.029) it makes undercutting unprofitable for the inelastic majority; as a public event it shifts beliefs, releasing the wave of coordinated increases. I estimate the model by simulated method of moments, and it reproduces the main war-to-cartel sequence and the principal fit moments, while leaving the tight late-rent cluster unmatched. I find the harm is a transfer to supra-competitive rents, with small deadweight loss because post-ban demand is inelastic.

*Keywords:* cartel formation; dynamic games; firm beliefs; loss-leader pricing; store traffic; advertising restrictions; retail pharmacy; equilibrium transition; simulated method of moments.

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

---

\*I thank the Chilean Competition Authority (FNE) for data access, and the editor and anonymous referees for invaluable comments. Financial support from the 2020 Gambling Award Fellowship is gratefully acknowledged. All errors are my own. An open-access version of this paper is available at <https://arxiv.org/abs/2512.22917>.

# 1 Introduction

Between December 2007 and April 2008 the three retail pharmacy chains that together account for 92% of pharmacy sales in Chile (Cruz Verde (CV), Farmacias Ahumada (FASA; FA in equations and tables) and Salcobrand (SB)) coordinated price increases on more than two hundred prescription and over-the-counter medicines. The episode was documented in detail by Chile's National Economic Prosecutor's Office (*Fiscalía Nacional Económica*, FNE) and the Competition Tribunal (*Tribunal de Defensa de la Libre Competencia*, TDLC) in two contemporaneous court records (the FNE *Requerimiento* of 9 December 2008 and the TDLC *Sentencia* N° 119/2012<sup>1</sup>) and is the setting I study here, following earlier work on collusion among retail pharmacy chains [Alé-Chilet, 2016, 2018].

Through 2006 the chains discounted these medicines only occasionally; a persistent below-cost price war set in from January 2007. In August 2007 Cruz Verde launched a comparative-price advertising campaign (*Desafío*) that intensified it; in September 2007 the advertising self-regulation council (CONAR) ruled the campaign out, and on 6 November 2007 a Civil Court precautionary order upheld the restriction. The first coordinated price increase that held followed within weeks, on 3 December 2007, and the coordinated period ran to the FNE's investigation notice of 31 March 2008.

Two facts organize the analysis. The court record documents the channel: manufacturers' sales reps (*visitadores*) carried a single coordinated price proposal, or batch call, to the three chains in the same week. The event panel adds timing.<sup>2</sup> It dates each attempt, separates failures from holds, and distinguishes the first margin-restoring increase from a second supra-competitive increase, which I call the rent. Failures are frequent before and after the cartel but rare during it; many drugs receive two increases; and the increases arrive in laboratory batches.

I study the transition from the loss-leader competition that preceded the coordinated period to the cartel that followed. The convictions establish collusive communication, and the lab channel shows how the chains selected the coordinated price. What remains is why that price became self-sustaining once the ban removed the incentive to undercut. Initiation has received less empirical attention than cartel sustainability [Green and Porter, 1984]; here the channel is explicit and the build-up is fast. Unlike the 2.5-year gasoline rollout in Byrne and de Roos [2019], this episode unfolds over four months around the 6 November 2007 Civil Court order banning Cruz Verde's comparative-price advertising. The wholesale-cost record lets me identify a per-period spillover in Chilean pesos (CLP)

---

<sup>1</sup>FNE, "Requerimiento contra Farmacias Cruz Verde, Farmacias Ahumada y Salcobrand," 9 December 2008, available at [https://www.fne.gob.cl/wp-content/uploads/2011/03/requ\\_0009\\_2008.pdf](https://www.fne.gob.cl/wp-content/uploads/2011/03/requ_0009_2008.pdf); TDLC, "Sentencia N° 119/2012," 31 January 2012, available at [https://www.fne.gob.cl/wp-content/uploads/2012/01/Sentencia\\_119\\_2012.pdf](https://www.fne.gob.cl/wp-content/uploads/2012/01/Sentencia_119_2012.pdf).

<sup>2</sup>§2 defines a coordinated increase as all three chains raising a given drug by at least 15% within a week, with the level held.

and compute welfare across the transition.

Two forces sustained the war, and the episode itself shows both are real. A chain prices below cost only because the shopper drawn in by a cheap headline drug fills the rest of her basket at a profit; that store-traffic value is what repays the lost margin, so the war's persistence is itself evidence that undercutting paid. The second force explains why the war did not end on its own. A coordinated increase holds only if rivals are expected to match it. Through the war, attempts to lead out of it reverted: rivals were not expected to match, so belief stayed low and self-fulfilling. The advertising ban moves both at once. By collapsing price sensitivity ( $\hat{\alpha}$  from 0.103 to 0.029) it makes a cut steal too little traffic to pay, so for the inelastic majority of drugs no chain wants to undercut the coordinated price. As a public, dated event it also lifts the shared belief that a raise will be matched, releasing the wave. This belief interpretation is grounded in the record: the FNE describes unilateral increases as vulnerable to demand diversion, and describes the agreement as expanding as its success was verified, with laboratories used to coordinate and monitor it [Fiscalía Nacional Económica (FNE), 2008, Tribunal de Defensa de la Libre Competencia (TDLC), 2012]. Neither change alone reproduces the path (§5).

The paper contributes in three ways.

First, it gives an estimable model of how a cartel begins. I model the move from the below-cost price war to coordinated supra-competitive pricing with one set of primitives, roll it forward from January 2007, and match the war, failed attempts, cartel wave, late rent, and larger-laboratories-first order.

Second, a public event can coordinate beliefs rather than only change payoffs. The advertising ban acts as a focal, dated signal that shifts the expectation a raise will be matched; it selects, rather than creates, the coordinated equilibrium. I embed this in a two-layer laboratory-to-drug dynamic game and estimate eleven parameters by simulated method of moments. The knockouts show that coordination needs both the undercut incentive to vanish and the belief to shift. Leadership is disciplined by loss cushions and within-market shares, using the case record and business mix for the ordering and dispersion scales for observed leader shares.

Third, one advertising policy facilitated collusion. The comparative-price ad ban collapsed the very price sensitivity that made loss-leading pay, so undercutting no longer repaid its margin for the inelastic majority and the loss-leader war gave way to coordination. A consumer-protection rule aimed at advertising thus had an unintended pro-collusion side effect. The resulting welfare cost is distributional, a transfer of about CLP 3.4bn relative to the below-cost war prices over the cartel window, with small deadweight loss because post-ban demand is inelastic.

**Related literature.** Most empirical work on collusion studies how it is sustained [Green and Porter, 1984, Harrington, 2018] or recovers collusive conduct from prices and quantities [Igami, 2017, Clark and Houde, 2014, Miller and Weinberg, 2017, Miller et al., 2021]; how

a cartel begins has drawn less attention. The mechanics of reaching coordination run through communication and observable price moves: Sugar Institute communication harmonising practices [Genesove and Mullin, 2001], asymmetric gasoline adjustments acting as cartel-holding transfers [Clark and Houde, 2013], and distributor-run hub-and-spoke coordination [Chaves and Duarte, 2025]. The closest study of initiation, Byrne and de Roos [2019], documents a two-and-a-half-year experimental rollout in retail gasoline; my episode is its short, regulatory-triggered counterpart, where one advertising restriction flips a static undercut incentive so coordination forms in four months without multi-year learning. Two studies by Alé-Chilet examine this same Chilean episode: Alé-Chilet [2016] on the within-cartel ordering and Alé-Chilet [2018] on Salcobrand’s leadership as Bayesian signalling.

The model’s primitive is the cross-category traffic a multi-product retailer wins by being cheapest: the loss-leader logic in which below-cost “leaders” draw shoppers whose baskets repay the lost margin [Lal and Matutes, 1994, Thomassen et al., 2017, Chen and Rey, 2012, DeGraba, 2003, Florez-Acosta and Herrera-Araujo, 2020, Rao and Syam, 2001]. There store traffic sustains below-cost pricing; I run it in reverse. That restricting price advertising raises prices is an old finding, for eyeglasses [Benham, 1972] and retail pharmacy [Cady, 1976, Milyo and Waldfogel, 1999, Sinkinson and Starc, 2019], as is the broader result that an advertising ban can soften competition [Eckard, 1991]; I supply the channel: the ban leaves the store-traffic value untouched and acts through consumers’ price sensitivity, whose estimated collapse shrinks the traffic a cut wins so loss-leading stops paying.

The episode also belongs to a literature on regulation that reshapes conduct in unintended ways [Carranza et al., 2015, Ryan, 2012, Cicala, 2015]. Here the consumer-protection advertising rule worked on two margins. It changed payoffs (collapsing the price sensitivity that made loss-leading pay, which made the coordinated price self-enforcing), and it shifted beliefs: a public, dated event that let the chains expect a coordinated increase to be matched. In this second role the rule acts like a focal point or commitment device, much as a price ceiling can become a focal point for tacit coordination [Knittel and Stango, 2003, Lewis, 2015]; the difference here is that the coordination is explicit and the rule released it rather than capping it. Closest is Alé-Chilet and Atal [2020]: a trade association lets many physicians coordinate the same low-to-high move; here the device is instead a public regulatory event, working through demand (a belief shift and the elasticity collapse) among only three chains.

Finally, firms often reach a new equilibrium by adaptive learning rather than jumping to it: fictitious play over three years in UK electricity [Doraszelski et al., 2018], learning-to-price after liquor privatisation [Huang et al., 2022], strategic ability in deregulation [Goldfarb and Xiao, 2011], biased beliefs in market power [Aguirregabiria and Jeon, 2020], menu-cost rigidity [Aguirregabiria, 1999, Kano, 2013]. My transition is the opposite limiting case: a dated event flips the undercut incentive and the belief I estimate is a fast

transient, so the wave is mechanical rather than learned. For the multi-product cartel structure I draw on Igami and Sugaya [2022] and for elasticity benchmarks on Grabowski and Vernon [1992]; methodologically I follow the simulated-method-of-moments tradition for dynamic games [McFadden, 1989, Pakes and Pollard, 1989, Bajari et al., 2007] and identify demand as in Berry [1994].

**Outline.** Section 2 describes the data and the regulatory record; Section 3 the descriptive patterns the model rationalises (below-cost pricing, the price-tier transition, and Salcobrand-led timing); Section 4 estimates nested-logit demand and the ban-induced collapse in price sensitivity, the pre-ban level on the competitive window and the post-ban level on the cartel-excluded weeks [Berry, 1994]; Section 5 sets out the structural model, shows it reproduces the war-to-cartel transition, and quantifies the welfare consequences; and Section 6 concludes.

## 2 Institutional Setting and Data

### Industry and regulatory record

The Chilean retail pharmacy market in 2007 was concentrated. Cruz Verde, FASA, and Salcobrand together held approximately 92% of nationwide pharmacy sales (FNE Requerimiento 2008, chap. II)<sup>3</sup>, and were estimated to account for as much as 95% of pharmacy revenue more broadly [Díaz and Galetovic, 2015]. Each chain set its list price for each medicine centrally at the corporate level, valid across all of that chain’s stores nationwide; the “coordination” I study in this article therefore refers to coordinated movements of the three chains’ centrally set national list prices, not to local store-level pricing.

Demand for these medicines is price-inelastic and physician-driven: the prescribing doctor, not the consumer, chooses the molecule, and a patient cannot substitute across branded products without a new prescription. The Competition Tribunal describes consumers as captive through the prescription (*cautivos a través de la receta médica*) [Tribunal de Defensa de la Libre Competencia (TDLC), 2012]. Generic competition is limited and largely outside the chains: only 22.5% of the molecules in my sample face a generic substitute, and those are sold mainly through independent pharmacies rather than the three chains, which stock predominantly branded products bought through common wholesale networks. This institutional inelasticity is the basis for the small quantity response, and hence the transfer-not-efficiency-loss welfare finding, of §5. It does not rule out substitution across the three chains for the same molecule: that store-choice margin is what comparative price advertising activated, and it is the margin estimated in the demand model of §4. This was not the first such episode: in 1993–94 the Comisión Resolutiva

---

<sup>3</sup>The 92% figure is from the FNE complaint, p. 7; the 95% figure from Díaz and Galetovic [2015] is a later academic estimate covering both prescription and OTC pharmacy revenue.

found the same three chains to have raised the prices of 80 ethical medicines in concert, a sanction the Corte Suprema affirmed [Comisión Resolutiva H. de Defensa de la Libre Competencia, 1995].

**Marginal-cost symmetry across the three chains.** The three chains face nearly identical wholesale-cost schedules: manufacturers publish a single suggested retail price (*precio de venta a público sugerido*, PVPS) per molecule, incorporating a 20–25% gross-margin band and sent to each chain alike, so wholesale-acquisition cost is common across chains for each drug. The one substantive asymmetry is Cruz Verde’s vertical integration with the distributor Socofar, which affects margin retained at the wholesale stage, not the marginal cost of acquiring the drug. I therefore treat the wholesale cost  $c_{jt}$  as common across chains, using Salcobrand’s disclosed wholesale-cost series (§2) as that common  $c_{jt}$ . Chain-level demand, by contrast, is not forced symmetric: the demand stage (§4) estimates a chain-by-drug intercept  $\phi_{ij}$  absorbing persistent differences in how a drug sells across the three chains, whereas the dynamic model adds a chain-specific loss-cushion asymmetry  $\psi_i$  that carries the leadership: the non-pharmacy business that absorbs part of the war loss, so the least-cushioned chain is most exposed and leads the escape (calibrated from the business mix, §5).

**Origin of the price war.** The loss-leader price war that this article analyses was triggered by an advertising campaign that Cruz Verde launched in August 2007 under the slogan “Cruz Verde Challenge: Low Prices Without Competition” (*Desafío Cruz Verde, Precios Bajos Sin Competencia*). The campaign published side-by-side price comparisons for selected products against FASA, claiming systematically lower prices.<sup>4</sup> All three chains then cut prices well below cost on the products that consumers were most likely to be familiar with, to defend pharmacy traffic share. This is the loss-leader equilibrium whose collapse drives the cartel-formation episode in this article. I measure the intensity of the war by the below-cost prevalence (Fig. 1, monthly revenue shortfall from below-cost selling), which rises tenfold from June to November 2007.

**The regulatory record.** Two regulatory events bracket the period of interest. The chronology below documents them; §5 takes up the roles. The first event is a sequence of three mutually reinforcing instruments that, between September and November 2007, terminated Cruz Verde’s comparative-advertising campaign and the loss-leader price war it had driven. On 7 **September** 2007, the board of the Advertising Self-Regulation Council (*Consejo de Autorregulación y Ética Publicitaria*, CONAR), acting on a complaint filed by FASA, issued a ruling in Case ROL 704/07 finding that Cruz Verde’s campaign infringed

---

<sup>4</sup>FNE, “Requerimiento contra Farmacias Cruz Verde, Farmacias Ahumada y Salcobrand,” 9 December 2008, chap. II § 3, identifies the price war (*guerra de precios*) as launched by Cruz Verde in August 2007, with FASA and Salcobrand reacting by cutting their own prices on the same products to defend share. The press coverage in *El Mercurio*, *Economía y Negocios* (late September / early October 2007, article ID 34637) and the *Diario Financiero* of 12 September 2007 documents the campaign and FASA’s reaction.

Articles 4, 6, 10 and 22 of the Chilean Code of Advertising Ethics (*Código Chileno de Ética Publicitaria*) and ordering Cruz Verde to cease the campaign.<sup>5</sup> CONAR, founded in 1987, is the self-regulatory organ of the Chilean National Advertisers Association (Asociación Nacional de Avisadores, ANDA); it adjudicates advertising-ethics complaints under that private code and its rulings are not judicially enforceable.<sup>6</sup> The board confirmed the ruling on reconsideration on 5 October 2007 and the CONAR ethics tribunal (*Tribunal de Ética*) confirmed it on appeal later in October. On 31 **October** 2007 the National Consumer Service (Servicio Nacional del Consumidor, SERNAC) added an administrative judicial complaint under the consumer-protection law.<sup>7</sup> The first coercive judicial instrument came on 6 **November** 2007, when the 17<sup>o</sup> Juzgado Civil de Santiago granted a *medida precautoria* requested by FASA in its parallel civil suit under Ley 20,169 (competencia desleal) and ordered Cruz Verde to discontinue the campaign, with the suit subsequently producing a definitive ruling on 2 July 2010 and a confirming sentence from the Court of Appeals of Santiago on 23 July 2012.<sup>8</sup> The first coordinated price increase in the event panel occurs 27 days later, on 3 **December** 2007.

The roughly thirteen-week gap from the CONAR ruling to the first coordinated increase is consistent with this dating. The non-binding CONAR ruling did not halt the war: below-cost selling intensified rather than abated through September–October 2007 (Figure 1). The binding 6 November Civil Court *medida precautoria* was the first instrument with the legal force to compel withdrawal. The four-week interval from it to the first coordinated increase on 3 December is short relative to the multi-week lab-mediated coordination cycle (single-manufacturer batches take weeks to organise across the three chains).

The cartel was brought to an end by the FNE's investigation that led to the formal complaint (*Requerimiento*) of 9 December 2008. The TDLC ultimately sanctioned all three chains, finding (i) that the chains had collusively coordinated price increases on 222 medicines, (ii) that the suggested retail margin on these medicines was 20–25%, and (iii) that the cartel's profits represented 2.8%/4.0%/3.1% of 2007 sales for FASA, Cruz Verde and Salcobrand respectively, with the 222 coordinated medicines accounting for roughly

---

<sup>5</sup>CONAR, Caso ROL 704/07 ("Farmacias Ahumada S.A. con Farmacias Cruz Verde S.A."), Directorio resolution of 7 September 2007; reconsideración of 5 October 2007; and Tribunal de Ética decision later that month. The compiled jurisprudence is at [https://www.conar.cl/wp-content/uploads/2011/04/Jurisprudencia\\_Conar-2007.doc](https://www.conar.cl/wp-content/uploads/2011/04/Jurisprudencia_Conar-2007.doc).

<sup>6</sup>The Sep–Oct CONAR rulings were morally binding but not coercive, which is why the price-war prevalence kept rising through October as Cruz Verde appealed; the binding instrument was the 6 Nov Civil Court injunction documented below.

<sup>7</sup>SERNAC press communication, "Sernac denuncia a Cruz Verde por su publicidad," 31 October 2007, available at <https://www.sernac.cl/portal/604/w3-article-918.html>.

<sup>8</sup>17<sup>o</sup> Juzgado Civil de Santiago, "Medida precautoria," 6 November 2007 (reported by CIPER Chile, "El dossier del caso farmacias," 9 April 2009, available at <https://www.ciperchile.cl/2009/04/09/el-dossier-del-caso-farmacias-asi-se-subieron-los-precios-segun-fasa/>); and definitive ruling of 2 July 2010, confirmed by the Court of Appeals of Santiago on 23 July 2012 (Rol 4155-2010).

16% of national pharmaceutical sales (the Farma segment, 2006–2009 average).<sup>9</sup> I treat these events as exogenous in the structural model, each with its own role: the September CONAR ruling ( $t_{\text{CONAR}}$ ) is the non-binding precursor that ends the comparative campaign; the binding November ban ( $t_{\text{ban}}$ ) both collapses demand (the price-sensitivity drop, §4) and, as a public signal, shifts the common belief (the focal jump, §5); and the March FNE notice ( $t_{\text{FNE}}$ , Oficio 419) raises enforcement on the rent tier (the post-Oficio scrutiny, the skepticism  $m$  the chains work off as they survive), not the coordination belief. The belief jump is dated to the binding November event, not the September ruling (§5).

**The batch-call coordination mechanism.** The coordination in this episode is explicit rather than tacit: pharmaceutical manufacturers in Chile distribute updated price recommendations through dedicated sales representatives (*visitadores*), who call on each chain’s central purchasing office on a periodic basis. Manufacturers publish a suggested retail price (*precio de venta a público sugerido*, PVPS) per drug (a list price that incorporates the standard 20–25% gross margin (§2)) and visitadores deliver revised PVPS schedules to each chain simultaneously. When a chain accepted a manufacturer’s new PVPS, it raised its prices on that manufacturer’s entire current portfolio within a single weekly window: this is a lab-batch call. The TDLC *Sentencia* 119/2012 documents that Salcobrand and rivals used the visitadores channel as a bilateral coordination conduit: Salcobrand would first commit to the manufacturer’s new PVPS, and the visitador would carry that commitment to Cruz Verde and FASA as the signal that triggered their follow-on raises.<sup>10</sup> Because the PVPS covers every presentation of a molecule from a single manufacturer, each visitador call naturally generates a batch across that lab’s full portfolio. Every batch call in November 2007 failed: the initiating chain raised prices and rivals declined to follow. The first sustained three-chain coordinations emerged in the week of 3–9 December (Online Table 4 lists the principal batches).

## Data

**Data.** Three sources, all from the public case record, underlie the analysis. **(i) Prices and quantities.** A transaction-level panel of daily list prices and unit sales for each of the 222 drugs named in Annex I of the 2008 FNE *Requerimiento*, at the three chains (Cruz Verde,

<sup>9</sup>Figures from the FNE *Requerimiento* (9 December 2008) and the TDLC *Sentencia* 119/2012 (both cited above with their URLs). The profit split in (iii) is the FNE’s quantification, reproduced in the *Sentencia*; the 16% figure in (iv) is the Tribunal’s own finding (Considerando 41), measured as the coordinated medicines’ share of national pharmaceutical sales.

<sup>10</sup>TDLC *Sentencia* 119/2012 (available at [https://www.fne.gob.cl/wp-content/uploads/2012/01/Sentencia\\_119\\_2012.pdf](https://www.fne.gob.cl/wp-content/uploads/2012/01/Sentencia_119_2012.pdf)), conduct findings; FNE *Requerimiento* 9 December 2008, chap. II, identifying the 222 coordinated medicines across 36 therapeutic categories (IMS classification). The FNE’s post-cartel market study Fiscalía Nacional Económica [2020] documents the continuing promotional role of visitadores; the FNE complaint and TDLC judgment document their role as the cartel conduit.

FASA, Salcobrand), obtained from the FNE under Chile’s transparency law (Ley No. 20.285), 2006–2008. Wholesale costs are Salcobrand’s, disclosed to the TDLC. I estimate demand on this daily panel; Online Appendix D gives the construction. **(ii) Event panel.** From the price series I build one event panel, combining a 15% coordination threshold with a more permissive 12% daily detector (a 12% move over the prior 42-day median, within a  $\pm 7$ -day rival-follow window).<sup>11</sup> It records price moves in both directions. A coordinated increase is all three chains raising a drug by at least 15% within a 7-day window with the level held, labeled a success, a failed attempt, or a unilateral hold by whether rivals follow. Where the automated detector is ambiguous, the initiating chain, the success-versus-failed-attempt label, and a small number of coordination successes are hand-coded from the case record (Online Appendix B gives the coding rules), so these structural moments are partly hand-curated rather than purely mechanical. A price decrease, or war deviation, is a chain cutting its price at least 15% below its own 2006 baseline for at least two weeks during the pre-cartel war; the panel records 814 such war deviations in total (the structural moment in §5 is a narrower windowed count, not this raw total). After coordination, a rarer defection is a chain cutting below both the cheaper rival and its own cartel plateau (51 events, about 16 of them deep and lasting past the March 2008 FNE notice). The depth of these below-cost decreases and the timing of the failed attempts are moments the structural estimation weights heavily. This panel is the source for the event counts in §3. Of the 222 drugs named in the TDLC *Sentencia*, 220 (99.1%) show at least one successful three-chain coordination (Online Appendix D details the mapping). **(iii) Documentary record.** The TDLC *Sentencia* No. 119/2012 (conviction on 206 drugs and fines near US\$38M), the FNE complaint (Requerimiento, Rol C No. 184-08), the CONAR self-regulatory ruling (704/07), the 6 November 2007 civil-court injunction, and FASA’s 2007 Humphreys credit-rating report [Humphreys Clasificadora de Riesgo, 2007]. All are publicly available with their URLs (Online Appendix I).

### 3 Descriptive Evidence

**Market structure.** Three national chains (Cruz Verde, FASA, and Salcobrand) together hold close to 90% of formal pharmacy sales. They compete head-to-head for the same branded molecules, priced off a common centralized national list and bought through shared wholesale networks. Two features of this structure shape every fact below. Demand is largely captive (§2), so the chains compete for which store a shopper visits, not whether she buys; this is what makes a below-cost loss leader a rational bid for store traffic. The

---

<sup>11</sup>Online Appendix B validates this panel against the FNE official cartel indicator: the two largely agree, with a median timing difference of  $-2$  days and 79% of drugs within a week. I use the event panel as the primary outcome because it isolates the coordination event from the trailing reaction periods that contaminate the official panel.

chains’ own accounts show the basket is large enough to matter: in the pre-cartel years non-pharmaceutical lines (cosmetics, personal care, baby products, beverages) made up about a quarter of Cruz Verde’s profit (24%) and revenue (23%), so a customer a cheap headline drug draws in fills a substantial basket. The store-traffic value  $\mu$  the model estimates is the profit on the marginal traffic a deviation wins (the deal-seeking switchers), so it sits above this average per-visit margin. And coordination, when it comes, runs through the manufacturers’ batch calls rather than chain-to-chain, so the unit that moves is the laboratory’s portfolio, not the single drug.

The regulatory timeline is set out in §2: the war that Cruz Verde’s August-2007 campaign drove, the CONAR ruling and 6 November Civil Court order that ended it, and the FNE investigation that closed the cartel. Read against that timeline, the facts describe a loss-leader equilibrium whose intensity was accelerating through 2007 toward a level that could not be sustained. They then describe a coordinated wave of price increases that restored the chains’ weighted pharmacy margin to roughly the historical 20–25% band. I establish three facts and organise the section around them: **(i)** pre-ban below-cost pricing whose intensity was accelerating, not stationary (§3); **(ii)** a sharp price-tier transition that restored the historical margin and was not a cost shock (§3); and **(iii)** Salcobrand-dominant leadership and a four-month wave organised laboratory-by-laboratory, the largest laboratories first (§3). Sections 4 to 5 build the structural model that rationalises the magnitudes; this section establishes that they are there to be explained.

### Fact 1: Below-cost pricing accelerated before the ban

The standard formal argument for a loss-leader equilibrium is that firms expect the per-period loss on the loss-leader to be recovered from contemporaneous spillovers on other categories (§5). A natural diagnostic of whether such an equilibrium is sustainable is whether the loss is approximately stationary or whether it grows over time. The below-cost gap is visible directly in Figure 1: across 2007 each chain’s quantity-weighted price (Panel B) runs below the wholesale cost (Panel A), and the gap widens through the year. Panel C aggregates, for each chain, the monthly revenue shortfall from below-cost sales: that is,  $\sum_{j,t} \max(c_j - p_{ijt}, 0) \cdot q_{ijt}$  summed across drugs and days within each month.<sup>12</sup>

Through 2006 and the first half of 2007, the monthly shortfall is roughly steady at ~CLP 50–100 million per chain. From mid-2007 onwards, every chain’s shortfall begins to grow, and the growth accelerates into November–December 2007, with Cruz Verde reaching a peak of roughly ~CLP 900 million per month. The chains were losing roughly

<sup>12</sup>The cost benchmark  $c_j$  is the drug’s wholesale cost from the Salcobrand cost series, which the data record from November 2007; for earlier months I hold  $c_j$  at its first observed level, because the wholesale price of these off-patent molecules is stable over the window. Because the pre-November benchmark is backward-extrapolated, the level of the early shortfall is measured with more error than the post-November level; the acceleration into late 2007 is robust to this, but its exact onset month should be read as approximate.

an order of magnitude more in late 2007 than they were a year earlier. Within weeks of the first coordinated increase, the shortfall is back to baseline.

This pattern is hard to reconcile with the loss-leader equilibrium having been “stable” through 2007. It motivates the structural measurement: the firms moved as the net payoff to staying put in the war deteriorated.

### **Fact 2: The price-tier transition**

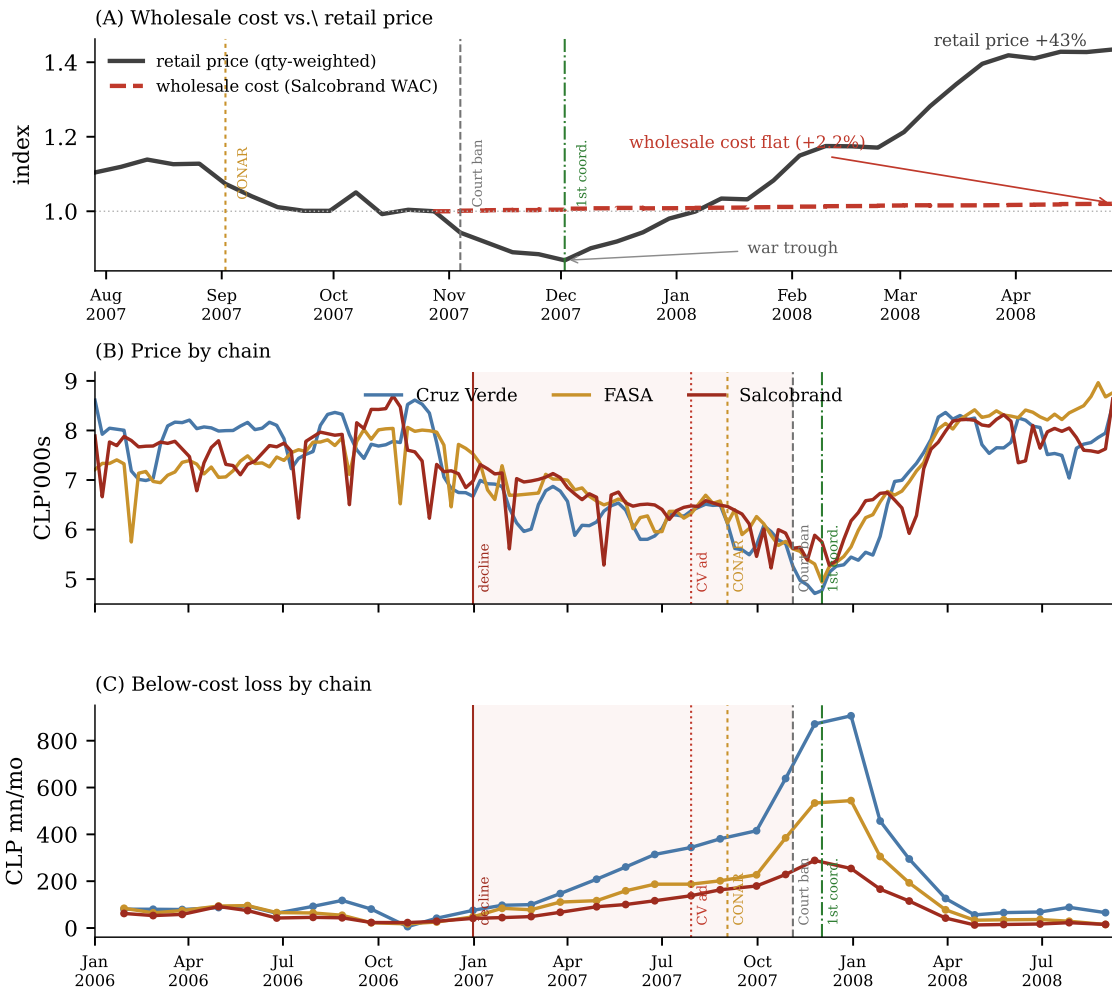
Figure 1 (Panel B) shows the weekly quantity-weighted average price for each chain over the 2006–2008 window. Three features stand out. First, prices co-move throughout: across the 156 weeks of data, the three chains are within roughly  $\pm$ CLP 500 of one another at the weekly frequency, consistent with the centralised national list-pricing described in §2. Second, prices drift downward from mid-2006 through late 2007 (from  $\sim$ CLP 7,000 to  $\sim$ CLP 5,000), with the dip steepening between the CONAR ruling and the first coordinated increase. Third, in a single sharp step in December 2007–February 2008, the chain-level price returns to and exceeds the 2006 level. This step-up was not a cost shock. Over the same coordination window the chains’ wholesale costs rose only a few percent, about 3.3% on Salcobrand’s WAC series disclosed to the Competition Tribunal (December 2007–May 2008; +2.2% over the longer November–May window of Figure 1). Retail prices over the same window rose 28–60% across therapeutic categories [Tribunal de Defensa de la Libre Competencia (TDLC), 2012]. A cost shock cannot account for a price path that fell during the war and then rose far more than costs during the recovery: the reversal reflects coordinated conduct, not input costs.

The coordinated step restores a sustainable, nonnegative margin. Quantity-weighted markups move from substantially negative ( $-12$  to  $-16\%$ ) at the war floor to small positives at the first ( $\ell = 0 \rightarrow 1$ ) step, still below the 20–25% suggested-retail band the TDLC documents (*Sentencia* 119/2012, Cap. XI). A second push ( $\ell = 1 \rightarrow 2$ ) overshoots to +34 to +41%, into the region the Tribunal flags as supra-competitive. This asymmetry maps onto the structural model directly: the first step is the regime flip (the ban turns the one-shot undercut gain  $G_j$  negative, restoring a sustainable margin), the second a selective rent-extraction step taken by only the most inelastic drugs (§5). Table 1 decomposes the markup by chain and tier.

### **Fact 3: Leadership and timing**

Salcobrand led 247 of the 363 coordination events (68%) and dominated the cartel wave (211 of 284 coordinations). This holds even though Cruz Verde bore the heaviest pre-ban below-cost exposure (31.3% of drug-weeks below cost against Salcobrand’s 25.6%, FASA between). The structural model (§5) resolves this through the loss cushion. Relative to

Figure 1: Wholesale cost, prices, and below-cost loss



**Note:** (A) Salcobrand wholesale-cost index (dashed) and quantity-weighted retail-price index (solid), both indexed to the cost-series start, November 2007–May 2008. (B) Weekly quantity-weighted average price by chain across the 222 coded drugs, 2006–2008. (C) Monthly below-cost revenue shortfall by chain,  $\sum_j \max(c_j - p_{ijt}, 0) q_{ijt}$ . The CONAR self-regulatory ruling (7 September 2007) and the binding Civil Court order (6 November) are marked in Panels B–C.

FASA’s international operations and Cruz Verde’s Socofar structure, Salcobrand appears to be the most exposed domestic pharmacy chain, so I treat it as having the least outside business to cushion the war’s below-cost losses. This makes the war costliest for it,

Table 1: Markups by tier and chain

	Tier 0 (pre-coord.)	Tier 1 (1st coord.)	Tier 2 (2nd coord.)
<b>Panel A. Unweighted mean markup</b>			
Cruz Verde	0.025	0.302	0.596
FASA	0.041	0.304	0.592
Salcobrand	0.047	0.310	0.610
<b>Panel B. Quantity-weighted mean markup</b>			
Cruz Verde	-0.158	0.105	0.371
FASA	-0.154	0.089	0.342
Salcobrand	-0.122	0.127	0.405
<b>Panel C. Count of products with negative markup</b>			
Cruz Verde	133	60	12
FASA	131	59	14
Salcobrand	127	53	11

**Note:** Panel B: quantity-weighted mean markup. Panel C: count of chain-drug-weeks with negative markup. Markup  $m_{ijt} = (p_{ijt} - c_{jt})/p_{ijt}$ . Tiers correspond to the coded coordination events (§2).  $N = 222$  drugs.

giving it the strongest incentive to lead the escape even though the larger Cruz Verde cut deeper. The external record supports the ordering: FASA had substantial Mexico/Peru operations; Cruz Verde was tied to the Socofar distribution structure; and Salcobrand was the domestic, pharmacy-focused chain undergoing a management change in 2007 [Humphreys Clasificadora de Riesgo, 2007, Fiscalía Nacional Económica, 2010, Tribunal de Defensa de la Libre Competencia, 2012, Diario Financiero, 2007]. The case record also confirms the observed order: the TDLC records Salcobrand as the usual first mover in the 1–2–3 mechanism, with rivals following within a few days [Tribunal de Defensa de la Libre Competencia (TDLC), 2012]. In the post-investigation rent the below-cost loss is gone and market share governs instead, so the largest chain Cruz Verde narrowly overtakes Salcobrand (36 initiations to 34), the SB → CV handoff the model reproduces. The loss-cushion ordering is read from the chains’ business mix, not from prices alone. Table 2 reports the full leadership split.

The wave ran 3 December 2007 to 31 March 2008: 216 of the 220 first coordinations (Tier 0 → 1) fall within these 19 weeks, whereas the rent steps (Tier 1 → 2) ratchet on through year-end. Over the full episode the chains made 220 first coordinations and 143 rent steps, the 363 events the structural model targets, 284 of them in the cartel-window success moment. Against these stand 194 failed or reverted attempts from January 2007 on (Panel B gives the full 2006–2008 ledger by window). The ordering was organized at the laboratory level: each manufacturer’s visitador delivered a batch across its portfolio,

Table 2: The coordination ledger: leadership, regime, and magnitude

<b>Panel A. Below-cost exposure and leadership, by chain</b>						
	CV	FASA	SB			
<i>Pre-ban below-cost exposure</i>						
Fraction below cost	31.3%	26.9%	25.6%			
Median ( $p - c$ ) ( $10^3$ CLP)	0.70	0.87	0.98			
<i>Leadership, by period (initiator)</i>						
War (pre-cartel)	0	0	2			
Cartel wave (Dec'07–Mar'08)	14	59	211			
Post-investigation rent	36	7	34			
All coordinations	<b>50</b>	<b>66</b>	<b>247</b>			
of which first-tier ( $0 \rightarrow 1$ )	13	42	165			
of which rent ( $1 \rightarrow 2$ )	37	24	82			
Share of all ( $n = 363$ )	13.8%	18.2%	68.0%			
<b>Panel B. Events by kind and regime window</b>						
Event kind	Pre-decline (2006)	War/dec. (Jan– 5Nov'07)	Pre-cartel (6Nov– 2Dec'07)	Cartel (3Dec– 31Mar)	Post-invest. (1Apr– Dec'08)	Total
First coordination (Tier $0 \rightarrow 1$ )	0	1	1	216	2	<b>220</b>
Rent step (Tier $1 \rightarrow 2$ )	0	0	0	77	66	<b>143</b>
Failed attempt (reverted)	86	99	23	4	67	279
Unilaterally held	62	50	12	20	42	186
<b>Total events</b>	<b>148</b>	<b>150</b>	<b>36</b>	<b>317</b>	<b>177</b>	<b>828</b>
<b>Panel C. Coordination steps by tier transition and magnitude</b>						
	Large ( $\geq 25\%$ )	Med. (15–25)	Small (12–15)		< 12%	
$0 \rightarrow 1$ (first)	177	31	3		9	
$1 \rightarrow 2$ (rent)	51	47	21		24	
<b>All coord.</b>	<b>228</b>	<b>78</b>	<b>24</b>		<b>33</b>	

Magnitude bins sum to 363 coordination steps (220 first + 143 rent).

**Note:** Chains are Cruz Verde (CV), FASA, and Salcobrand (SB). Panel A decomposes the 363 coordination events by chain. Below-cost exposure: drug-weeks with  $p_{ijt} < c_{jt}$  in the pre-ban window (weeks 1–94; 9,945 drug-weeks per chain). Leadership: the initiating chain of each event (the 220 first coordinations  $0 \rightarrow 1$  and the 143 rent steps  $1 \rightarrow 2$ ), split by the war (pre-cartel), cartel (3 Dec 2007–31 Mar 2008, to the FNE *Oficio*) and post-investigation periods, so the rows reconcile to the full ledger ( $2 + 284 + 77 = 363$ ). Salcobrand leads the war and the cartel wave; Cruz Verde overtakes it in the post-investigation rent (the SB  $\rightarrow$  CV handoff). Panel B sub-windows: “Pre-decline” = 2006 (prices flat at  $\approx 7.7$ ); “War/decline” = Jan 2007–5 Nov 2007 (the decline begins January 2007 and Cruz Verde’s August-2007 comparative-price campaign intensified it); “Pre-cartel” = 6 Nov–2 Dec 2007 (civil-court injunction to the day before the first three-chain success); “Cartel” = 3 Dec 2007–31 Mar 2008 (to the day the FNE opened its investigation, notice No. 419); “Post-investigation” = 1 Apr–Dec 2008. The 220 first coordinations form the cartel in the Dec–Mar window. The rent then deepens through 2008: 77 of the 143 rent steps fall in the cartel window, the other 66 post-investigation, as the regime price climbs toward 15. Panel C magnitude is the price increase over the prior tier (the war price for the  $0 \rightarrow 1$  step, the Tier-1 price for the  $1 \rightarrow 2$  rent): “large”  $\geq 25\%$ , “medium” 15–25%, “small” 12–15%, the rest smaller week-aligned steps. Most are large (228 of 363 steps are  $\geq 25\%$ ). Coordinations and rents are from the dynamic ledger (which tracks the rent past the panel’s April cutoff); failed lifts and unilateral holds are the coded event series.

so timing is set by which laboratory engaged first. Larger laboratories coordinated first: laboratory membership, not a drug’s own size, elasticity, or below-cost depth, sets the order (Table 3). Before the cartel held, the chains were already testing: 35 attempts in the 6 November to 2 December window, 23 failed or reverted and 12 unilaterally held (Figure 4).

Table 3: Lab membership, not drug characteristics, sets the coordination order

Predictors of the first-coordination day (OLS, $N = 219$ )	$R^2$
Drug covariates (size, elasticity, below-cost gap)	0.03
+ drug-type fixed effects	0.09
+ <b>manufacturer (laboratory) fixed effects</b>	<b>0.61</b>


**Note:** Each row regresses the calendar day of a drug’s first coordinated increase on the listed predictors (the  $J = 220$  coordinated drugs,  $N = 219$  after dropping one with incomplete pre-ban characteristics; continuous regressors standardised, larger = later). The manufacturer fixed effect is decisive: adding it raises  $R^2$  from 0.09 to **0.61**. By contrast, the drug’s own size, elasticity, and below-cost depth explain only 0.03 of the timing variance. The order is set by which laboratory engaged, not by the drug. The laboratories with the highest pre-ban trading volume coordinate first (Figure 2). Standard errors clustered by manufacturer ( $G = 37$ ).

Figure 2 shows the size gradient directly. A laboratory’s median first-coordination week falls with its size at a correlation of  $-0.29$ , where size is the log of its total pre-ban revenue, the same stake  $W_L$  the structural model assigns it. A larger laboratory plausibly moves first because it carries more SKUs to bring under the coordinated umbrella and so makes a more focal coordination device: starting with its batch synchronises the most drugs at once. This negative correlation does work in the estimation. The moment  $\text{corr}(\log \text{ lab size, week}) = -0.29$  disciplines the lab-engagement parameters  $\lambda$  and  $\tau_e$  in the structural model (§5, Table 8), so the larger-labs-first ordering is matched by the model rather than assumed. Table 4 records the principal batch calls: every batch call in October and November failed, and the first sustained three-chain coordinations emerged in the week of 3–9 December.

Figure 3 illustrates one episode. The weekly price for Sildenafil Citrate (Drug #80) drifts down through mid-2007, then steps up within a single week in early December 2007 (tier 0  $\rightarrow$  1, led by Salcobrand) and again in early April 2008 (tier 1  $\rightarrow$  2); the three chains co-move tightly across both transitions. The official binary indicator  $\text{col}_{ijt}$  records only the first step (it is monotone per drug), so it understates the cartel’s reach by about one-third. It records fewer steps than the event coding in 127 of the 137 multi-round drugs (Online Appendix B), which is why I use the tier ladder throughout.


The failure rate (Figure 4) traces the regime. It runs  $\approx 97\%$  through the first four weeks of November, below 10% by mid-December, near zero through March. Then it snaps back up after the FNE issues Oficio N° 419 on 31 March 2008, when new coordination stops while the cartel holds and the rent deepens. The reverted lifts in this last window

Figure 2: Larger laboratories coordinate first



**Note:** Each point is one laboratory (37 labs, 221 coordinated drugs), at its size and the median week of its drugs' first three-chain coordination. Lab size is the laboratory's total pre-ban revenue (log CLP), summed across its drugs, the stake  $W_L$  the structural model assigns it. Larger labs coordinate earlier (correlation  $-0.29$ ; the structural model reproduces this at  $-0.38$ , Table 8). Point area is proportional to the laboratory's number of drugs. Dashed and dash-dotted lines mark the binding 6 November court ban and the first sustained coordination, 3 December 2007.

Figure 3: Example coordination: Sildenafil price path



**Note:** Weekly chain price for Sildenafil Citrate (Drug #80) with the coded tier transitions marked. The three chains move together within the same week at each tier step; Salcobrand initiates both.

are failed new attempts on not-yet-coordinated drugs, not defections: the 220 members

Table 4: Early lab-batch calls failed; December calls held

Dates	Laboratory	Drugs	Initiator	Outcome
Oct 7–13	Andromaco	3	Salcobrand	Failed
Nov 8–13	Laboratorio Chile	19	Salcobrand	Failed
Nov 30–Dec 8	Andromaco	6	Salcobrand, FASA	Partial (2/6; the two holds FASA-led)
Dec 7–11	Bago	4	Salcobrand, FASA	Success
Dec 12–15	Tecnofarma	4	FASA	Success
Dec 18–20	Saval	5	Salcobrand, FASA, Cruz Verde	Success

**Note:** Principal lab-batch calls, October–December 2007, from the event panel. A “batch call” is a span of about a week in which a chain raised prices on several drugs from one laboratory at once; Drugs counts the distinct drugs (DrugIDs) in the batch and Initiator lists each leading chain with  $\geq 1$  drug (ordered by drug count). Outcome is the three-chain result: “Failed” = the initiating chain reverted with no rival follow, “Success” = all drugs held three-chain, “Partial” = a fraction held. The two early calls (Andromaco, October; Laboratorio Chile, November) failed. The first sustained three-chain coordinations came in December: the Andromaco call of late November began holding on 4 December (two drugs, FASA-led), and the Bago, Tecnofarma, and Saval calls held in full across the next two weeks. The lab-level size gradient (larger laboratories coordinate first) is shown in Figure 2.

largely hold their coordinated price<sup>13</sup> (Table 2, Panels B–C, gives the full ledger by regime window and magnitude). Salcobrand leads the failed re-coordination push (38 of the 46 post-investigation attempts, 83%), the same war-exposed chain that led the escape from it.

The dynamic structural model of §5 takes this loss-leader regime as the period-0 state and asks who exits it and when; the next three sections build that model.

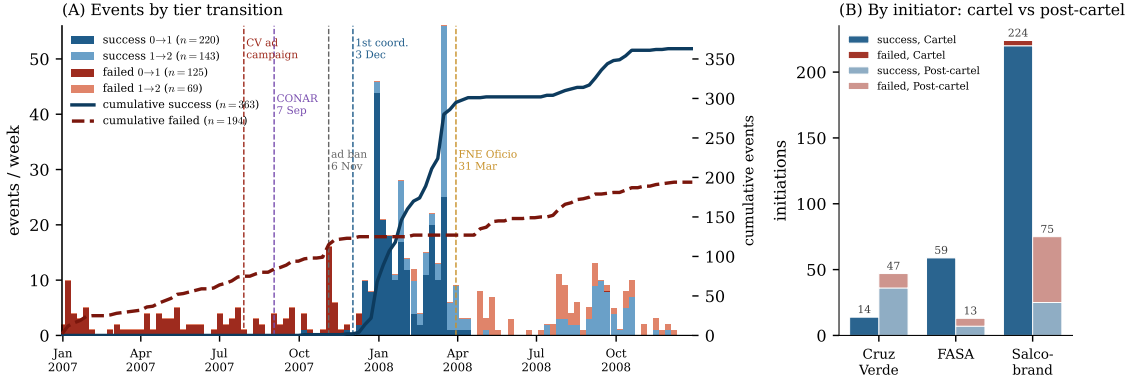
## 4 Demand

### Model

A consumer in week  $t$  decides whether to fill a prescription for drug  $j \in \{1, \dots, J\}$  at one of the three national chains  $i \in \mathcal{I} = \{CV, FA, SB\}$ , or to use the outside option (an independent pharmacy or no purchase). I use a two-level nested logit in which the three

<sup>13</sup>A small number of coordinated drugs defect later: about sixteen show a deep, lasting cut by one chain after the March 2008 FNE Oficio, the chain falling below the held price without rejoining. These cuts are isolated (the other two chains hold) and fall where undercutting is unprofitable ( $G_j^{01} < 0$ ). They read as enforcement-era retreat rather than competitive undercutting, and their small number matches the model’s prediction that post-ban defection is rare.

Figure 4: Coordination events over the episode



**Note:** Coordination events over January 2007–December 2008, decomposed two ways. Over the episode the ledger records 363 successful coordination steps (220 first-tier 0→1, 143 rent 1→2) against 194 failed or reverted attempts (125 at 0→1, 69 at 1→2). **Panel A:** weekly events split four ways, successful 0→1 (dark blue) and 1→2 (light blue) versus failed 0→1 (dark red) and 1→2 (light red); the right axis plots cumulative successes (solid) and failures (dashed). (Failed attempts carry no recorded tier; I assign each to 0→1 or 1→2 by whether the drug had already coordinated by the attempt date.) **Panel B:** initiations by chain, each split successful versus failed and cartel-window versus post-investigation. Vertical dashed lines (chronological): Cruz Verde’s August comparative-price campaign, the 7 September CONAR ruling, the 6 November ban, the 3 December first coordination, and the 31 March FNE official notice No. 419.

chains form a single inside nest for each drug and the outside option is the only competing nest.<sup>14</sup> Consumer  $h$ ’s indirect utility from buying drug  $j$  at chain  $i$  is

$$U_{hijt} = u_{ijt} + v_{hjt} + (1 - \sigma) \varepsilon_{hijt}, \quad \sigma \in [0, 1), \quad (1)$$

where  $v_{hjt}$  is the nest-level taste shock common to the three chains,  $\varepsilon_{hijt}$  is an i.i.d. Type I extreme-value shock, and the nesting parameter  $\sigma$  measures within-nest correlation across the three chain alternatives. Let  $D_t \equiv \mathbf{1}\{t \geq t_{\text{CONAR}}\}$  be the demand-regime indicator. The demand equation uses September as the empirical split because the CONAR ruling opens the restricted-advertising run-up in the data; this is not the structural payoff node. The structural model dates the binding collapse in price sensitivity, with the belief shift it brings, to the November court ban that made the restriction enforceable (§2). Mean utility is

$$u_{ijt} = X_{jt}\beta - (\alpha_0 + \alpha_1 D_t)p_{ijt} + \phi_{ij} + \xi_{ijt}, \quad \alpha_0 > 0, \quad \alpha_0 + \alpha_1 > 0, \quad (2)$$

<sup>14</sup>Drug  $j$  is defined at the molecule-dosage level; the three chains carry the same brand-package portfolio for each  $j$  (Section 2). Prescription pre-selects  $j$ , so substitution across drugs is absorbed in the outside good.

where  $p_{ijt}$  is the transaction price (CLP 000s),  $\phi_{ij}$  is a chain-by-drug fixed effect,  $X_{jt}\beta$  collects time-varying drug characteristics, and  $\xi_{ijt}$  is an unobserved demand shock. The price coefficient is decomposed into a pre-ban level  $\alpha_0$  and a post-ban shift  $\alpha_1$ ; the null of no ban effect on price sensitivity is  $H_0: \alpha_1 = 0$ . The sign convention  $\alpha_0, \alpha_0 + \alpha_1 > 0$  makes (2) consistent with downward-sloping demand on both sides of the ban.<sup>15</sup>

Aggregating over consumers yields the within-nest share  $s_{iljt}$ , the nest share  $s_{jt} = \sum_{i \in \mathcal{I}} s_{iljt}$ , and the outside share  $s_{0jt} = 1 - s_{jt}$ . Taking logs and differencing relative to the outside good [Berry, 1994] gives the single pooled estimating equation

$$\underbrace{\ln s_{ijt} - \ln s_{0jt}}_{Y_{ijt}} = X_{jt}\beta - \alpha_0 p_{ijt} - \alpha_1 (D_t \cdot p_{ijt}) + \sigma \ln s_{iljt} + \phi_{ij} + \xi_{ijt}. \quad (3)$$

I estimate equation (3) in two passes: pre-ban weekly GMM, and a post-ban re-estimation on the cartel-excluded weeks. The weekly panel is the natural prescription-refill unit and the standard frequency in pharmaceutical demand studies [Ellison et al., 1997, Björnerstedt and Verboven, 2016]. The per-drug coefficients  $\hat{\alpha}_j$  instead use the daily panel, whose far larger count of within-drug, cross-chain price observations gives the stronger first stage a drug-specific slope needs.

The market size for drug  $j$  is

$$N_j^{\text{wk}} = 7 \max_t Q_j^{\text{wk}} + 1, \quad (4)$$

where  $Q_j^{\text{wk}} = q_{CV,jt} + q_{FA,jt} + q_{SB,jt}$  is the total weekly three-chain quantity. The factor of seven makes the potential market larger than the observed peak refill week, so the outside option includes eligible non-buyers as well as consumers using independent pharmacies. The constant +1 prevents a degenerate  $s_{0jt}$  at the peak observation; the chain-by-drug fixed effect  $\phi_{ij}$  absorbs any remaining level mis-scaling.<sup>16</sup>

The baseline specification imposes a single  $\sigma$  and a single pre-ban price coefficient  $\alpha_0$ . With three chains per drug-week, in-sample variation does not separately identify drug-specific price coefficients from  $\phi_{ij}$ , so the cross-product dispersion in implied elasticities reported in Section 4 reflects variation in prices and shares, not preference heterogeneity. I relax the homogeneity restriction in Section 4 below.

<sup>15</sup>Throughout, I report estimates of  $\alpha_0$  and  $\alpha_0 + \alpha_1$  as the relevant price sensitivities;  $\alpha_1$  is the ban-induced shift.

<sup>16</sup>Online Appendix C reports market-size sensitivity checks.

## Identification and Instruments

Two regressors in (3) are endogenous: price  $p_{ijt}$  (correlated with  $\xi_{ijt}$  through unobserved demand shocks) and the within-nest share  $\ln s_{ijjt}$  (correlated with  $\xi_{ijt}$  mechanically). The instrument set is four variables:

1. **Wholesale cost (drug average)**,  $\bar{c}_j$ : the weekly average of Salcobrand's recorded wholesale cost for drug  $j$ , averaged across weeks. Drug-level cost variation pins down the cross-section of  $\alpha$ .
2. **Wholesale cost (weekly)**,  $c_{jt}$ : Salcobrand's wholesale cost in week  $t$ . Within-drug, week-to-week cost variation pins down the time-series identification of  $\alpha$ .
3. **Rival price 1**,  $p_{-1,jt}$ , the contemporaneous price of one rival chain.
4. **Rival price 2**,  $p_{-2,jt}$ , the price of the other rival chain.

The two cost instruments shift the marginal cost faced by all three chains symmetrically (the wholesale market is national), so they enter  $p_{ijt}$  but, conditional on  $\phi_{ij}$  and the time controls in  $X_{jt}$ , are excluded from  $\xi_{ijt}$ . The two rival prices follow the Hausman [1996] logic: they shift the within-nest share  $\ln s_{ijjt}$  through demand substitution but are otherwise excluded.

**Instrument validity across the ban.** The drug-average cost shifter and the rival-price instruments enter the headline pre-ban estimate. The weekly cost series begins in November 2007, so cost-only time-series identification is available only post-ban; the cost-only row in Table 5 is therefore a post-ban check. The two rival-price (Hausman) instruments are valid only when the chains compete: once they coordinate, rival prices co-move with the demand shock  $\xi_{ijt}$  and are no longer excluded. Pooling all ads-banned weeks, the over-identifying  $J$  rejects ( $J = 14.5$ ,  $p = .001$ ). The rejection is confined to the collusive cartel window (3 Dec 2007–11 Apr 2008). Dropping those weeks and re-estimating on the competitive ads-banned weeks, the same four instruments pass the over-identifying test ( $J = 4.4$ ,  $p = .11$ ; Table 6, note ‡). I therefore identify  $\alpha_0$  on the pre-ban weeks and the post-ban price sensitivity on the cartel-excluded competitive weeks, with  $\sigma$  pooled at its pre-ban value 0.393 so that price is the only endogenous regressor.

**The collapse does not hinge on the instruments.** The two endogenous regressors are not on the same footing. The within-nest share  $\ln s_{ijjt}$  is mechanically tied to  $\xi_{ijt}$  through the Berry inversion, so every nested-logit estimate instruments it; an OLS that does not returns an inadmissible  $\hat{\sigma} > 1$ . The own price is the contested case. To isolate it I hold  $\sigma$  at its estimated value and re-estimate the price coefficient under OLS, the rival-price instruments, and the wholesale-cost instrument. Table 5 reports the three. The price coefficients are close and all imply a large post-ban decline. Pre-ban the coefficient is 0.093 under OLS against 0.103 under the rival IV; post-ban it is 0.032, 0.029, and 0.030 under OLS,

the rival IV, and the cost IV. The collapse is 0.35 under OLS and 0.29 under the rival IV. The own-price endogeneity is therefore not large enough to drive the result, and the price collapse that drives the structural mechanism does not rest on the rival-price exclusion. Because the wholesale-cost series begins in November 2007, it identifies the post-ban coefficient; pre-ban identification comes from rival prices, and the OLS comparison shows the result is not sensitive to that choice.

Table 5: The price collapse is robust to the identification strategy

	OLS	Rival-price IV	Cost IV
$\hat{\alpha}_0$ (pre-ban)	0.093	0.103	—
$\hat{\alpha}^{\text{post}}$ (post-ban)	0.032	0.029	0.030
Collapse ratio	0.35	0.29	—

**Note:** Own-price coefficient,  $\sigma$  fixed at 0.393, under OLS, the rival-price (Hausman) IVs, and the wholesale-cost IV. The cost series begins November 2007, so the cost IV identifies only the post-ban coefficient. Collapse ratio =  $\hat{\alpha}^{\text{post}}/\hat{\alpha}_0$ .

## Estimation and Test Statistics

I estimate (3) by two-step IV-GMM [Berry, 1994, Conlon and Gortmaker, 2020] on the pooled chain-drug-week panel, absorbing the chain-by-drug effect  $\phi_{ij}$  with a within-group transform (instruments demeaned likewise). The moment condition is  $E[Z_{ijt} \xi_{ijt}(\theta)] = 0$  with  $\theta = (\alpha_0, \alpha_1, \sigma, \beta)$ ; each instrument enters at both its direct and ban-interacted level  $D_t \cdot Z_{ijt}$ , which separates  $\alpha_0$  from  $\alpha_1$  (§4). Standard errors treat  $\xi_{ijt}$  as heteroskedastic across  $(i, j, t)$ , matching the heteroskedasticity-robust GMM weight matrix.

On each window I report the first-stage  $F$ -statistics for  $p_{ijt}$  and  $\ln s_{iljt}$  [Stock and Yogo, 2005], the Hansen [1982] overidentification  $J$ , and the fitted-observed correlation in  $Y_{ijt}$  (the nested-logit analogue of  $R^2$ ). The  $J$  has  $q - k = 1$  degree of freedom on the pre-ban window, where  $\sigma$  is jointly estimated. It has  $q - k = 2$  on the post-ban windows, where  $\sigma$  is fixed. I impose no bounds on  $\alpha$  or  $\sigma$ .

## Estimates

Table 6 reports the demand estimates. The headline row is the pooled Model B result ( $\hat{\alpha}_0 = 0.103$ ,  $\hat{\sigma} = 0.393$ ; cartel-excluded post-ban  $\hat{\alpha}^{\text{post}} = 0.029$ ) that the structural model consumes; the lower block reports the per-drug coefficients  $\hat{\alpha}_j$  that give the model its drug-level heterogeneity. The headline pooled row (Panel A) uses the weekly panel and the weekly market-size calibration (4); the per-drug coefficients (Panel B) use the daily panel ( $T = 1,096$  days, 1 January 2006–31 December 2008, 222 drugs) with the daily

analogue  $N_j = 7 \cdot \max_t Q_j^{\text{daily}} / \hat{\rho} + 1$  (the factor of seven converts peak-day quantity to a weekly-equivalent market, and  $\hat{\rho} = 0.92$  is the three-chain share of Chilean pharmacy sales).

**Per-drug coefficients.** Per-drug  $\hat{\alpha}_j$  are 2SLS slopes on the cartel-excluded daily panel ( $\hat{\sigma}$  fixed at 0.393; rival-price IVs pre-ban, + Salcobrand cost post-ban); the structural model scales each pre-ban  $\hat{\alpha}_j$  by the pooled ratio 0.29. Online Appendix A.1 gives the per-drug specification.

Table 6: Heterogeneous price coefficient by drug class

**Panel A. Pooled headline** (homogeneous  $\alpha, \sigma$  pooled at the pre-ban value)

	$\hat{\alpha}$	$\sigma$	$N$	$J(p)$	
Pre-ban (cost + rival IV) SE	0.103 (0.005)	0.393 (0.107)	62,985	0.60 (.44)	identifies $\alpha_0, \sigma$
Post-ban, all weeks (incl. cartel) SE	0.022 (0.001)	[0.393]	41,769	14.5 (.001)	not used <sup>‡</sup>
Post-ban, cartel-excluded SE	<b>0.029</b> (0.001)	[0.393]	29,172	4.4 (.11)	used; $\alpha^{\text{post}}/\alpha_0 = 0.29$

**Panel B. Per-drug  $\hat{\alpha}_j$  by drug type:** 2SLS, Firm FE, cartel weeks dropped, pooled  $\hat{\sigma} = 0.393$ ; pre-ban rival-price IV, post-ban Salcobrand cost *and* rival IVs (daily panel)

Drug type	$n_{\text{drugs}}$	Pre-ban $\hat{\alpha}_j$	Post-ban $\hat{\alpha}_j$	post/pre
Chronic × Rx	130	0.123	0.023	0.19
Chronic × OTC	66	0.127	0.061	0.48
Acute × OTC	11	0.229	0.026	0.12
Acute × Rx	15	0.214	n.i.	—
All 222 drugs	222	0.130	0.029	0.22

<sup>‡</sup> On all ads-banned weeks the  $J$  test rejects ( $J = 14.5, p = .001$ ): cartel weeks are collusive, so rival prices fail exclusion. Dropping them restores validity ( $J = 4.4, p = .11$ ) and gives the post-ban  $\hat{\alpha} = 0.029$  ( $\alpha^{\text{post}}/\alpha_0 = 0.29$ ).

*Per-drug.* Each  $\hat{\alpha}_j$  is a 2SLS slope (rival-price IV pre-ban, + Salcobrand cost IV post-ban; cartel-excluded),  $\hat{\sigma}$  pooled at 0.393; the model scales each pre-ban  $\hat{\alpha}_j$  by the pooled collapse 0.29 (Panel A): the structural model thus uses the per-drug pre-ban level and the economy-wide collapse. The all-drug per-drug ratio in Panel B is 0.22 rather than 0.29 because the pooled pre-ban coefficient  $\hat{\alpha}_0 = 0.103$  lies below the per-drug median 0.130; the model uses the precisely-estimated pooled collapse, not the noisier per-drug posts.<sup>17</sup>

The per-drug coefficients line up with prior knowledge of pharmaceutical demand on two points. First, the acute-over-chronic ranking holds (refill prescriptions create switching costs that dampen price response). Second, the per-drug estimates collapse post-ban to the same all-drug 0.029, matching the pooled headline. The cost-only IV (0.030) and the

post-investigation-only window (0.028) reproduce it without using the rival prices at all.

## Goodness of Fit and Implied Elasticities

Three diagnostics describe how well  $\hat{\theta} = (\hat{\alpha}_0, \hat{\alpha}_1, \hat{\sigma})$  fits the data: within-nest share fit, the cross-sectional ranking of implied own-price elasticities against pharmaceutical benchmarks [Goldman et al., 2004], and a tier-2 holdout. The formulas are collected in Online Appendix A.

**(A) Within-nest share fit.** The structural within-nest share prediction (1) uses only the estimated parameters and prices (the observed share never enters the right-hand side). It is therefore a genuine prediction, not a mechanical contraction-mapping  $R^2$  [Conlon and Gortmaker, 2020, p. 32]. Predicted and observed within-nest shares correlate 0.73 (pre-ban), 0.86 (post-ban), and 0.78 pooled (Table 7).

**(B) Implied elasticities, benchmarked.** I evaluate the standard nested-logit formulas (2)–(3) at each drug’s pre-ban prices and shares. The median own-price elasticity is  $-1.75$ , with 99.5% of drugs negatively signed and all cross elasticities positive. Post-ban the median falls to  $-0.41$  at the collapsed price sensitivity and the coordinated prices, with 87% of drugs now inelastic. The cross-class ranking agrees in sign with the U.S. Medicare estimates of Goldman et al. [2004]: chronic essential categories price-inelastic, discretionary ones price-elastic. The classes a chain pushes below cost are the inelastic chronic refills, not the most elastic items (Spearman  $\rho = -0.62$  across 12 ATC1 classes; Online Appendix Table 1).

**(C) Tier-2 holdout.** Refitting on tier- $\{0, 1\}$  weeks and predicting  $\hat{s}_{i|jt}$  on the held-out tier-2 weeks via (1) (with  $\hat{\phi}_{ij}$  recomputed on the training subset) tests whether demand extrapolates to the fully-coordinated regime in which the welfare results of §5 are evaluated; the held-out correlation is 0.85.

Table 7: Demand goodness-of-fit diagnostics

	Pre-ban	Post-ban	Pooled
(A) Within-nest share fit (1), non-circular prediction			
Corr( $\hat{s}_{i jt}, s_{i jt}$ )	0.73	0.86	0.78
(C) Tier-2 holdout: predict on held-out tier-2 weeks			
Corr (tier-2 holdout)		0.85	

**Note:** Diagnostic (A) reports the correlation between predicted and observed within-nest shares from (1) on the pooled chain-drug-week panel; diagnostic (B) (category ranking) is in Online Appendix Table 1; diagnostic (C) reports the held-out tier-2 correlation.

A regime decomposition (Online Appendix A) shows the residual fit gap is the static

logit's week-by-week below-cost selection, immaterial for the tier-level welfare calculation of §5, which uses tier-level prices where shares are stable.

## 5 Dynamic Structural Model

This section embeds the store-traffic value and the followers' belief in a dynamic pricing game. The estimated collapse in price sensitivity at the advertising ban, with no change in costs, patience, or any other fundamental, flips the static incentive: undercutting pays in the loss-leader war and stops paying after the ban. The belief shift then selects the coordinated path, so the same model, rolled forward, reproduces the price war, the post-ban cartel wave, and the subsequent rent extraction. The game is non-stationary. The advertising ban, the February summer holiday, and a finite detection horizon all fall inside the sample window, so I solve a rolling-horizon weekly recursion, closed by a stationary post-ban terminal value, rather than a stationary infinite-horizon recursion. Beyond the estimated demand, the model adds a store-traffic value  $\mu$  and a small set of organizational, belief, and timing parameters governing how the coordination wave rolls out.

The descriptive evidence establishes the sequence: a below-cost war, failed attempts, a December cartel wave, and a later rent. It cannot, by itself, say how much of that sequence is due to the demand collapse, how much to the focal belief shift, or how large the counterfactual welfare cost is. The dynamic model is used for that narrower purpose. It weighs the mechanisms under one set of primitives, shuts them off one at a time, and carries the resulting price path into the welfare calculation.

### The dynamic game

Coordination resolves on two nested timing layers, the structure the case record forces. A laboratory engages at its own rate, and within an engaged laboratory each drug's lift is then decided. Larger laboratories engage first, so the documented order in which they coordinate is an outcome of the model, not an assumption.

The dynamic game has five pieces: the state (a); the per-period payoff and the static-Nash incentive flip that drives the war-to-cartel transition (b); the leader probability (c); the follower probability (d); and the belief through which the dated events enter (e).

**Solution concept.** The chains play a Markov equilibrium in subjective beliefs, not a rational-expectations Markov-perfect equilibrium (MPE). Given the market-wide belief  $b_t$  that a lift will be matched, each chain best-responds using the forward-simulated continuation values (9). The belief is not imposed to equal the equilibrium match probability; it is a Beta-Binomial posterior updated from matched and reverted lifts (14). Along the path it stays consistent with realized play, in the spirit of experience-based equilibrium [Fershtman and Pakes, 2012] and learning toward coordination [Doraszelski et al., 2018]. The war

is a low-belief trap, and the transition is a public signal that jumps the belief; a rational-expectations MPE assumes away exactly those dynamics. Given the belief path, (9) is a contraction (bounded payoffs,  $\delta < 1$ ), so the week-by-week roll-forward is well-defined.

**(a) State.** The model is built drug by drug. Each of the  $J = 220$  drugs is sold by the three chains, has chain-specific demand quality  $\hat{\phi}_{ij}$ , belongs to a laboratory  $L(j)$ , and carries its own war, restored, and rent prices  $(p_j^0, p_j^1, p_j^2)$ , cost  $c_j$ , and price coefficient  $\alpha_j$ , which together determine the one-shot undercut gain  $G_j(\alpha_t)$ . Tier prices and cost are read from the event ledger:  $p_j^0$  from below-cost war deviations,  $p_j^1$  from first coordinations,  $p_j^2$  from rent coordinations (imputed as  $1.12 p_j^1$  when absent), and  $c_j$  from implied costs across events. Their cross-sectional distribution, not a representative drug, enters the model. Each lab carries stake  $W_L = \sum_{j:L(j)=L} R_j$ . The weekly state is the tier vector  $\ell_t \in \{0, 1, 2\}^J$  (0 war, 1 restored, 2 rent), the advertising regime  $\alpha_t$ , and the calendar week.

**(b) Payoff and the static-Nash flip.** Drug  $j$ 's within-period demand is the estimated nested logit of §4: chain  $i$ 's share is

$$s_{ij}(\mathbf{p}_j; \alpha_j) = \frac{e_{ij}}{D_j} \frac{D_j^\omega}{1 + D_j^\omega}, \quad e_{ij} = \exp\left(\frac{\hat{\phi}_{ij} - \alpha_j p_{ij}}{\omega}\right), \quad D_j = \sum_i e_{ij}, \quad \omega = 1 - \hat{\delta}, \quad (5)$$

with  $\hat{\delta} = 0.393$  and the three chains in one nest. Each chain carries its own estimated quality  $\hat{\phi}_{ij}$ , so the chains split the nest asymmetrically (Cruz Verde 0.45, FASA 0.31, Salcobrand 0.24 at the coordinated price, matching the data); the coordinated-window intercepts predict the war- and rent-period within-shares out of sample to a median 1.6 percentage points (Online Appendix A.1). The only object the ban changes is the price coefficient: each drug keeps its own pre-ban elasticity  $\hat{\alpha}_j$ , and the ban scales every one of them down by the same pooled factor  $\alpha^{\text{post}}/\alpha_0 = 0.029/0.103$ . The per-drug post-ban interactions are individually too noisy to identify, whereas the pre-ban  $\hat{\alpha}_j$  that carry the cross-drug heterogeneity are cleanly estimated. The comparative-price ads that told shoppers who was cheapest are gone, so shares respond much less to price. Leadership, by contrast, is carried by the loss-cushion asymmetry  $\psi_i$ , so the least insulated chain, Salcobrand, leads the escape (§(c), §5).

Beyond the drug margin, each customer is worth a store-traffic value  $\mu$ , the chain's variable profit on the rest of her basket, so the per-period payoff is

$$\Pi_{ij}(\mathbf{p}_j; \alpha_j) = M_j s_{ij}(\mathbf{p}_j; \alpha_j) (p_{ij} - c_j + \mu), \quad (6)$$

with  $M_j$  the market size and  $c_j$  the wholesale cost. Store traffic is why pricing below cost ( $p_{ij} < c_j$ ) can be rational, the lost drug margin repaid by basket profit on the extra traffic [Lal and Matutes, 1994, Thomassen et al., 2017]; held fixed across the ban,  $\mu$  lets the ban's entire effect run through  $\alpha$ .

The transition turns on the traffic a deviation wins. A chain undercutting the coordinated price  $p^c$  to the war floor  $p^w$  gains share  $\Delta s = 0.188$  with ads on ( $\alpha = 0.103$ ) but only 0.050 with ads banned ( $\alpha = 0.029$ ), a 3.8-fold collapse. The cut pays only if the store-traffic profit on the stolen share clears the drug margin it gives up:

$$G(\alpha) = \underbrace{\mu \Delta s}_{\text{traffic profit}} + \underbrace{[s^{\text{dev}}(p^w - c) - s^{\text{coord}}(p^c - c)]}_{\text{drug-margin change} < 0} > 0 \iff \mu > \mu^*(\alpha) = \frac{-(\text{drug-margin change})}{\Delta s}. \quad (7)$$

Because  $\Delta s$  collapses 3.8-fold while the margin barely moves, the threshold jumps about 3.5-fold, from  $\mu_{\text{ads on}}^* = 8.6$  to  $\mu_{\text{ads banned}}^* = 30.3$ . The estimate  $\hat{\mu} = 9.26$  (thousands of CLP, anticipating §5) lies between the two thresholds,  $8.6 < 9.26 < 30.3$ : above the ads-on threshold, so the below-cost cut paid while ads ran, and below the ads-banned threshold, so the same cut stops paying once they are banned. Online Appendix G gives the analogous lone-leader share-loss distribution and its money decomposition. The follower faces the mirror choice, to match the lift or to harvest by staying cheap while the leader is expensive. Once the post-ban demand collapse makes the lift viable, matching beats harvesting for the inelastic majority; pre-ban, low belief and profitable undercutting make attempted lifts revert. Operating through demand alone, the ban thus flips the deviation incentive ( $G < 0$ ) into a one-shot best response for 168 of 220 drugs, with no appeal to patience or off-path discipline. The most-elastic 52 are held by the common-knowledge belief. (I test the observed deviations, not a global static best response.)

**From daily demand to a weekly payoff.** I estimate demand on the daily panel (§4), but the coordination game steps weekly. Chain  $i$ 's weekly payoff on drug  $j$  at tier  $\ell$  therefore scales the daily flow profit (6) up to a week,

$$\Pi_{ij}^\ell = 7 M_j^{\text{day}} s_{ij}^\ell (p_j^\ell - c_j + \mu) = N_j s_{ij}^\ell (p_j^\ell - c_j + \mu), \quad (8)$$

with  $M_j^{\text{day}}$  the daily market size and  $s_{ij}^\ell = s_{ij}(p^\ell; \alpha_{j,t})$  chain  $i$ 's estimated tier- $\ell$  share under the prevailing regime, and  $N_j \equiv 7 M_j^{\text{day}}$  the weekly-equivalent market.

**(c)–(d) Value, leadership, and the followers' constraint.** The value function carries the equilibrium tier transitions that are the substance of the model: a coordinated price is undercut back into the war, the war is escaped, and a restored price is deepened into the rent. I compute the forward-simulated continuation value, discounting the expected next-period value under the model's own one-step transition probabilities  $P_j^{x \rightarrow x'}(t)$  (lift, hold, undercut, deepen, revert) at the realized state,

$$V_j^x(t) = \Pi_j^x(\alpha_t) + \delta \sum_{x'} P_j^{x \rightarrow x'}(t) V_j^{x'}(t+1), \quad x \in \{0, 1, 2\}, \quad (9)$$

closed by the stationary post-ban perpetuity ( $\delta = 0.95$ ). These transition probabilities are not free parameters; they are the model's own choices, defined below:  $P^{0 \rightarrow 1}$  is the lift  $p_{ij}^\uparrow$  (12),  $P^{1 \rightarrow 2}$  the rent push (15), the reversions  $1 \rightarrow 0$  and  $2 \rightarrow 1$  the undercut  $p_{ij}^\downarrow$  (13), and  $P^{x \rightarrow x}$  the complement. I solve value and choice jointly to the fixed point: the value sets the escape gains  $\Delta V$ , and  $\Delta V$  in turn sets the choices that generate  $P$ . Write  $\Delta V_j^1 = V_j^1 - V_j^0$  for the escape value and  $\Delta V_j^2 = V_j^2 - V_j^1$  for the rent value; firms escaping the war do not foresee the rent, so  $\Delta V_j^1$  excludes it.

Whenever a chain is below cost, in a loss state, the loss cushion  $\psi_i$  governs. The least insulated chain bears the most loss, so it is the most desperate to act. I write  $\psi_i \in [0, 1]$  for chain  $i$ 's loss cushion, calibrated outside the model from its business mix. Farmacias Ahumada, with half its revenue and two thirds of its cash flow in Mexico, is the most insulated ( $\psi_{FA} \approx 0.8$ ). Cruz Verde, cushioned by the Socofar distribution structure documented in the FNE record, is middling ( $\psi_{CV} \approx 0.6$ ). Salcobrand appears most exposed and domestic/pharmacy-focused, so I set the exposed corner ( $\psi_{SB} \approx 0$ ).<sup>18</sup> The outside business cushions a fraction  $\psi_i$  of the below-cost war loss, so  $\psi$  enters through the payoff: chain  $i$ 's war-state flow profit is

$$\Pi_{ij}^0(\alpha_t) = \psi_i \Pi_j^1(\alpha_t) + (1 - \psi_i) \Pi_j^0(\alpha_t), \quad (10)$$

a convex combination that leaves a fully-exposed chain ( $\psi_i = 0$ ) bearing the entire loss  $\Pi_j^0$  and moves a fully-insulated chain's war payoff toward the coordinated level  $\Pi_j^1$ . In the choice rule I map the common continuation gain into an exposure-scaled escape value,  $\Delta V_{ij}^1 \equiv (1 - \psi_i) \Delta V_j^1$ . This is how the calibrated loss cushion enters the leader and follower choices; it is not a separate identity imposed on the value recursion (9). Under Type-I extreme-value shocks of scale  $\tau$ , chain  $i$  then leads a move out of a loss state with probability

$$\pi_{ij}(t) = \frac{\exp(\Delta V_{ij}^1(t)/\tau)}{\sum_k \exp(\Delta V_{kj}^1(t)/\tau)}. \quad (11)$$

The same exposure-scaled value enters the follower's incentive constraint below, so  $\psi$  is not confined to the leader. But the escape value swamps the one-shot undercut gain by an order of magnitude (Figure 5), so even the most-insulated chain's constraint stays slack: every chain still holds, and  $\psi$  moves who initiates without moving who follows. The leadership is thus carried by the loss cushion and the coordination count by the symmetric

---

<sup>18</sup>The calibration uses the ordering in external records, not the leadership shares themselves: Humphreys' 2007 FASA rating report documents the Mexico/Peru mix; the FNE market investigation and TDLC record document Socofar/Cruz Verde's distribution structure; and the TDLC/press record documents Salcobrand's 2007 ownership and management change.

store traffic.

Two events are loss states, and Salcobrand leads both. The below-cost war is one, so Salcobrand leads the escape, the ordering the business-mix  $\psi$  predicts. In the data Salcobrand initiates 74% of cartel-wave increases, which the model reproduces (69%). A failed rent push is the other, because a push the followers reject reverts the drug below cost, and Salcobrand leads those failures too, four-fifths of the post-ban failed attempts. Once a price holds above cost the loss is gone and  $\psi_i$  no longer functions. There the leader is the largest chain, Cruz Verde, selected by market share, the SB  $\rightarrow$  CV handoff that holds the durable rent.

A lift survives only if the followers match and no rival then undercuts. A follower's belief  $b_t$  is exactly that probability, so it holds rather than undercuts iff  $b_t \delta \Delta V_{ij}^1 > G_j^{01}(\alpha_t)$ , where  $G_j^{01}$  is the static gain to undercutting a coordinated price (7), positive while ads run and negative once the ban flips it. Because each of the two followers must expect the other to match, the belief enters squared, and a lift survives with conditional probability

$$p_{ij}^{\uparrow}(t) = \pi_{ij}(t) b_t^2 \text{logit}^{-1}\left(\left[b_t \delta \Delta V_{ij}^1 - G_j^{01}(\alpha_t)\right]/\tau\right). \quad (12)$$

The mirror down-move undercuts a held coordination back into the war. The slack  $s_{ij} = b_t \delta \Delta V_{ij}^1 - G_j^{01}$  is the belief-weighted continuation value the chain forgoes minus the one-shot undercut gain (positive when the chain prefers to hold, negative when undercutting pays), and the most-tempted chain undercuts in week  $t$  with probability

$$p_{ij}^{\downarrow}(t) = \lambda_{\text{cut}} \text{logit}^{-1}(-s_{ij}/\tau), \quad (13)$$

the mirror of the lift (12). Whether a coordinated price is undercut is decided by the most-negative slack, and an undercut returns the drug to the war floor ( $1 \rightarrow 0$ ), reversing the coordination the lift built. When undercutting pays ( $s_{ij} < 0$ , the gain exceeds the continuation) the probability rises toward  $\lambda_{\text{cut}}$ ; when it does not ( $s_{ij} > 0$ ) it falls toward zero. A rent reverts  $2 \rightarrow 1$  by the same rule on its own slack. The cartel's stability is therefore an outcome of the ban flipping  $G_j^{01}$  negative, not an assumed no-defection: defection is available every period but unprofitable. The rate  $\lambda_{\text{cut}}$  is itself a free scalar, estimated with the others: the model micro-founds when a profitable undercut is available (the slack gate), not its baseline intensity. The deterrence is the slack gate, not the fitted rate: perturbing  $\lambda_{\text{cut}}$  by  $\pm 50\%$  barely moves the coordination count over that range and mainly rescales the war's deviation intensity (Online Appendix G.4).

**Relation to punishment-supported collusion.** The slack gate is the analogue of a no-deviation constraint, but for this transition model rather than for a repeated game with off-path punishments. A chain undercuts when the gain  $G$  exceeds the belief-weighted

continuation value  $b_t \delta \Delta V$  it forgoes. The penalty is endogenous and mild: the drug drops one tier and re-enters the estimated transition, so the continuation is read from the forward-simulated value recursion [Igami and Sugaya, 2022], not an assumed off-path route. For the coordinated price the ban makes  $G_j^{01} < 0$ , so holding it is a one-shot best response and no off-path discipline is needed. The supra-competitive rent is different:  $G_j^{12} > 0$ , so the rent sticks only while  $b_t^{\text{rent}} \delta \Delta V_j^2 > G_j^{12}$ . The model leaves out an off-equilibrium discipline phase by design; in the data post-investigation defections are sparse, isolated retreats rather than the start of a retaliatory war. This weak rent discipline captures the rent's incentive and Cruz Verde leadership but not the synchronized late-2008 surge (§5).

**Timing within the week.** The moves resolve in a fixed order. A leader is drawn for each not-yet-coordinated drug (11), the two followers simultaneously decide whether to match, and the lift succeeds only if both do (12). The most-tempted chain may then undercut a held coordination at rate  $\lambda_{\text{cut}}$ , and the belief  $b_t$  updates from the week's matched and reverted lifts (14) before the next week begins.

**(e) Belief and the dated events.** Coordination is dynamically sustainable in both regimes (§5). What holds the chains in the war is not that the coordinated price is unsustainable but that no one expects it to be reached, so the transition selects the coordinated equilibrium rather than creating it. The case record has the same logic in institutional form: a lone raise sends demand to the rivals, and the FNE describes the agreement as extending as its success was verified, with laboratories coordinating and monitoring the agreement [Fiscalía Nacional Económica (FNE), 2008, Tribunal de Defensa de la Libre Competencia (TDLC), 2012]. I model that uncertainty as a belief about whether a lift will be matched. The belief is a single, market-wide Beta–Binomial posterior over a shared prior  $(a_0, b_0)$ , updated by Bayesian learning from the running tally of matched ( $S_t$ ) and reverted ( $F_t$ ) lifts pooled across all 220 drugs, so a coordination on one drug raises the same belief that gates every other (the cross-product channel that carries the wave). The **binding November ban** (week  $t_{\text{ban}} = 96$ ) is commonly observed, so every firm performs the same update, adding  $\zeta \approx 307$  pseudo-successes at the ban; the rent step reads the same belief under a fragility discount  $e^{-\chi} \approx 0.12$ . The three are one object,

$$b_t = \frac{a_0 + S_t + \zeta \mathbf{1}\{t \geq t_{\text{ban}}\}}{a_0 + b_0 + S_t + F_t + \zeta \mathbf{1}\{t \geq t_{\text{ban}}\}}, \quad b_t^{\text{rent}} = e^{-\chi} b_t, \quad (14)$$

the war posterior before the ban ( $t < t_{\text{ban}}$ ), the coordination posterior after it (a shift, not a reset, since the war's failures stay in  $F_t$ ), and the discounted rent posterior. The estimated prior is low ( $a_0/(a_0+b_0) \approx 0.01$ ): almost no one expects a lift to be matched, the few tried are undercut, and the pessimism is self-fulfilling. The ban lifts the belief in one step to about one-half, past the squared-belief threshold; it then climbs only to about seven-tenths over the wave and never near one, so the post-ban dynamics are governed by the one-

time incentive flip rather than by belief continuing to rise. The belief enters the lift (12) squared (both followers must expect a follow) and the rent once (the rent only deepens a coordination the chains already hold). The **September CONAR ruling** is a non-binding precursor; the **December onset**  $t_0 \approx 103$  is endogenous, the lag the laboratory reach-out generates; the **March FNE Oficio** raises enforcement on the rent alone (below). The rent hazard thus tracks the wave at a lower level and opens slowly. A rent-eligible drug (one whose own demand is inelastic enough that the higher price still pays,  $\Delta V_j^2 > 0$ ) is pushed  $1 \rightarrow 2$  in week  $t$  with probability

$$p_j^2(t) = b_t^{\text{rent}} r_t, \quad (15)$$

the discounted rent belief (entering once, not squared: the rent only deepens a coordination the chains already hold jointly) times a learned enforcement belief  $r_t$ .<sup>19</sup> The Oficio sets off a contested retreat in the rent. Right after the notice  $r_t$  is low, so most rent pushes are rejected and revert, and these failures persist rather than completing the rent. The chains also suspend new pushes for the  $\varphi \approx 9$ -week peak-scrutiny pause, holding the rent without risking the coordination tally. As quiet weeks accumulate  $r_t$  recovers, the surviving pushes begin to stick, and the rent climbs to its late peak.<sup>20</sup> The pause window  $\varphi \approx 9$  weeks is read from the institutional timeline (the Oficio, 31 March, plus the chains' observation lag), not estimated; the skepticism  $m$  is the one estimated enforcement parameter ( $t_0$  a fixed date).

Three objects discipline the rent. The discounted rent belief  $b_t^{\text{rent}}$  scales the probability of a push, the eligibility  $\Delta V_j^2 > 0$  determines which drugs can profitably deepen, and the learned enforcement belief  $r_t$  governs the post-Oficio pause and recovery. The pre-ban rent feasibility is therefore demand-driven rather than coordination-belief-driven. Forcing the coordination belief to one over-produces the  $0 \rightarrow 1$  wave but leaves the rent untouched before the ban: a further increase never pays while demand is still elastic, so no drug deepens until the ban makes demand inelastic, regardless of what the chains believe (§5). The rent push (15) therefore carries no static-best-response gate of its own; the eligibility already screens out drugs that would statically defect.

The same primitive (12) runs the pre-ban war: while ads are on  $G_j^{01} > 0$ , so any coordinated lift is undercut and reverts. Before the laboratories organize ( $t < t_0$ ) each not-yet-coordinated drug is tested directly at the follower belief itself (the leader raises

---

<sup>19</sup> $r_t$  is the chains' posterior that a rent push is safe from the regulator, not a calendar trend. After the Oficio they do not know whether the FNE will sanction the rent, and they treat each week survived without a sanction as evidence the threat is receding. This is a Beta-Binomial survival posterior  $r_t = (1 + k_t)/(1 + k_t + m)$ , where  $k_t = t - t_0$  is the weeks elapsed since the Oficio and  $m$  is the skepticism the notice instills. It starts low,  $r_{t_0} = 1/(1 + m)$ , and recovers as quiet weeks accumulate. No success term enters, so the belief cannot amplify itself.

<sup>20</sup>A one-time negative belief shock at the Oficio, the natural alternative, does not reproduce this: it depresses the rent hazard uniformly, so the rent flattens below the observed plateau, never accelerates, and generates no persistent failures; I verified this directly. The learned belief instead suppresses early matches, sustains the contested failures, and recovers only as the cartel survives, producing the late climb.

only if it expects a follow),

$$\Pr\{\text{drug } j \text{ tested for a lift in week } t\} = b_t [1 + g \mathbf{1}\{t \geq \text{Aug 2007}\}], \quad t < t_0, \quad (16)$$

the most war-exposed chain (Salcobrand, lowest  $\psi_i$ ) leading and Cruz Verde's August Desafío campaign intensifying by  $g$ ; almost every test is undercut, the reverts concentrating where  $G_j^{01}$  is largest, until the November ban flips  $G_j^{01} < 0$  and the war collapses.

**Two-layer organization.** The cartel was organized lab-by-lab through the manufacturers, the intermediaries the prosecutor identifies.<sup>21</sup> Each week a chain reaches out to a not-yet-organized lab  $L$  when the coordination it would unlock is worth more than the traffic it would forgo,

$$\Pr\{\text{chain } i \text{ organizes lab } L \text{ in week } t\} = \lambda \left[ \frac{1}{\tau_e} (b_t \delta \Delta V_{iL}^1(t) - G_L^{01}) \right]_0^1 \text{hol}(t), \quad (17)$$

summing the escape value and undercut gain over  $L$ 's drugs. Two ordering facts emerge from this reach-out rather than being imposed: larger labs (a larger summed surplus) cross the threshold first, so the documented larger-labs-first sequence is an outcome; and the most war-exposed chain, Salcobrand, with the largest escape value reaches out first and leads. Engagement is one-time and bounded (active for a spell of  $w$  weeks, then disengaged); conditional on it, each drug is pushed up at its own belief-gated rate (12), drugs coordinating at staggered weeks, and a coordinated drug is not profitably undercut once the ban flips  $G_j^{01}$  negative, so the cartel once formed holds.

## Estimation and fit

**The parameters.** The store-traffic value enters as a single common  $\mu$ . Leadership is calibrated with external discipline. In a loss state the leader is set by the loss cushion  $\psi_i$ , taken from each chain's business mix. In a profit state it is set by observed market share. Two dispersion scales then turn those externally grounded orderings into leader shares. The dynamic parameters are the common store-traffic value  $\mu$  and the reach-out scale and temperature  $(\lambda, \tau_e)$  governing laboratory engagement (17) and its onset  $t_0$ . The belief block adds the shared belief prior  $(a_0, b_0)$  together with the ad-ban pseudo-successes  $\zeta$  and the rent-tier signal penalty  $\chi$  that together micro-found the two belief shifters. The remaining parameters are the post-Oficio enforcement skepticism  $m$  that the learned rent-match belief works off, the symmetric down-cut rate  $\lambda_{\text{cut}}$ , and the campaign strength  $g$ . Fixed are the weekly discount  $\delta = 0.95$ , the Type-I EV choice scale  $\tau$ , the holiday factor  $\text{hol}$ , and the engagement window  $w = 26$  weeks, calibrated to the cartel's roughly six-

<sup>21</sup>"utilizaron a los laboratorios para coordinar y monitorear el acuerdo" [Fiscalía Nacional Económica (FNE), 2008].

month documented active span rather than estimated. The per-drug objects that drive the dynamics (each drug’s undercut and rent gains  $G_j^{01}$ ,  $G_j^{12}$  and shares, computed once from its own  $\alpha_j$ , prices, and cost) are inputs, not estimated here.

**Forward simulation.** For a candidate  $\theta$  I first build the per-drug value functions from the rolling-horizon recursion (9), then roll the weekly state  $\ell_t$  forward from January 2007 (all drugs in the war,  $\ell_0 \equiv \mathbf{0}$ ) over the 105-week window through December 2008. Week numbers are inherited from the full panel beginning January 2006, so January 2007 is week 52, the window ends at week 156, the binding November ban falls at week 96, and the coordination onset  $t_0 \approx 103$  in late December 2007, about seven weeks later. Each week resolves in order: while  $t < 96$  the below-cost war cheating (16), its failed attempts the only events; the ban at week 96 both collapses price sensitivity and, as a public event, jumps the coordination belief; and from the endogenous onset  $t_0 \approx 103$  on (once the laboratories have organized, the lag the reach-out generates) laboratory engagement (17), the belief-gated  $0 \rightarrow 1$  coordinate step (which succeeds only if both followers match and every chain holds its incentive constraint, else a failed attempt that reverts), and the  $1 \rightarrow 2$  rent step for the drugs whose own demand makes it pay (post-Oficio, an unmatched rent push reverts, a coordinated drug otherwise rarely defecting).  $B$  replications under common random numbers are averaged into the moment vector  $\widehat{m}(\theta)$ ,<sup>22</sup> and I minimize the weighted distance to the data moments,

$$\hat{\theta} = \arg \min_{\theta} \sum_k \left( \frac{\widehat{m}_k(\theta) - m_k^{\text{data}}}{w_k} \right)^2, \quad (18)$$

[McFadden, 1989, Pakes and Pollard, 1989, Bajari et al., 2007] by a derivative-free search (Powell), chosen because the objective is a step function of the discrete tier transitions. The criterion is multi-modal, so I search from several starts and keep the global optimum, which carries strictly lower loss than the nearest competing mode at every simulation size. Search and simulation use fixed seeds, so  $\hat{\theta}$  is reproducible. The hyperparameters I set rather than estimate are the replication count  $B$ , the moment weights  $w_k$ , the number of starts and the per-start iteration budget, the random seeds, and the three fixed values ( $\delta, \tau, \text{hol}$ ). Because the  $0 \rightarrow 1$  wave-path checkpoints are largely redundant with the onset, peak, and spread moments, I down-weight them so repeated path moments do not dominate the criterion.

**Standard errors.** The step-function criterion has no well-defined moment Jacobian, so the standard sandwich variance does not apply. I instead report a nonparametric bootstrap standard error: I resample the 220 drugs with replacement and re-estimate the model on each of 300 replications, each warm-started at the estimate, and take the

---

<sup>22</sup>The criterion stacks 38 moments against the 11 estimated parameters, so the model is over-identified by 27 restrictions; the moments are grouped in Online Appendix Table 4.

standard deviation across the draws (Table 8). The one parameter a symmetric standard error does not summarize cleanly is the ad-ban jump  $\zeta$ , which is identified only from below. No resample places  $\zeta$  under about 180, so the data sharply require a large focal jump but cannot bound it from above: once the jump is large enough to flip beliefs the wave saturates, and larger values are observationally equivalent (Online Appendix E.1). I read  $\zeta$  as a lower bound,  $\zeta \gtrsim 180$ , not a point estimate with a symmetric standard error.

**Identification.** Each parameter maps to a distinct moment (Online Appendix Table 4). Two cross-sectional moments discipline the micro-foundations. The who-cheats-most moment, the rank correlation between a drug’s pre-ban cheat frequency and its undercut gain  $G_j^{01}$ , identifies the benefit-ordered war (16). The who-takes-rent moment, the  $\alpha_j$  gap between drugs that reach the rent tier and those that stop, identifies the endogenous  $1 \rightarrow 2$  selectivity. The store-traffic mean  $\mu$  scales every payoff and rests on no single moment, though the below-cost war floor is the most informative and bounds it from below<sup>23</sup>; the belief prior  $(a_0, b_0)$  and the ad-ban signal  $\zeta$  are read from the regime split of the failed attempts,  $\chi$  from the smaller and later  $1 \rightarrow 2$  count,  $m$  from the post-investigation failure rebound,  $(\lambda, \tau_e, t_0)$  from the onset, peak and lab-size $\times$ timing gradient, and  $g$  from the August intensification.

**Estimates.** The eleven estimated parameters are in Table 8 (Panel A); the table carries the cell-by-cell fit, and I read off only the interpretive content. The store-traffic value is now a single common  $\hat{\mu} = 9.26$ , sitting just inside the ads-on threshold above which undercutting pays; the chain asymmetry has moved out of  $\mu$  and into the calibrated loss cushion  $\psi_i$  and market share. Cruz Verde’s category financials support the loss-leader economics: it sold chronic drugs below cost (chronic-drug profit  $-28$  M CLP/month) while harvesting acute (+2,299) and non-pharma (+1,088 M CLP/month) sales off the resulting traffic, a basket many times the drug loss. These aggregates are not used to estimate  $\mu$ , because they mix drug-driven and autonomous purchases.<sup>24</sup> The data contain no basket-level complement sales, so  $\mu$  is identified only from the chains’ own below-cost pricing; its level is bounded by that revealed behavior rather than measured against observed cross-category margins.

The model coordinates 214 of the 220 drugs and reaches 141 of the 143 rent steps: the 168 inelastic drugs hold on their own because undercutting is unprofitable, and the focal belief shift carries another 46 of the most-elastic 52. Coordination on a drug holds only if it is incentive-compatible for every chain, so the binding constraint is whichever chain

<sup>23</sup>The complement basket whose profit  $\mu$  captures is not in the case record, so  $\mu$  is identified from the chains’ revealed willingness to price below cost (the war floor), not measured from non-drug sales; its magnitude is that revealed-preference inference, an identifying assumption on the spillover rather than a direct measurement.

<sup>24</sup>Outside the case record, the *International Directory of Company Histories* (Farmacias Ahumada S.A.) records the same loss-leader model: dispensing was “a money-loser . . . compensated with high earnings on other products, including . . . health-and-beauty aids.”

is most tempted to undercut (the smallest per-chain slack). A drug fails to coordinate when even one chain's undercut gain  $G_{ij}^{01}$  stays positive. That is why the binding object is the product-firm incentive compatibility of (5), not a single product-level constraint. All three leadership moments line up under this leadership calibration. In the loss states  $\psi_i$  governs: Salcobrand leads 69% of the cartel wave (data 74%) and 91% of the post-early failed attempts (data 83%). In the profit state market share governs: Cruz Verde leads 68% of the post-investigation rent (data 65%), the SB→CV handoff (Table 8). These three leader shares are calibration targets; the external primitives pin each phase's leader but not the full ordering, so the runner-up positions are noisier (in the cartel wave, for instance, the model places Cruz Verde second where the data place FASA). The whole fit comes at an SMM loss of 16.59.

Table 8: Estimated parameters and model fit


<b>Panel A. Estimated parameters (SMM, eleven parameters)</b>		
Parameter	Est.	(SE)
$\mu$ store-traffic value (common)	9.26	(0.18)
$\lambda$ reach-out scale	0.63	(0.14)
$\tau_e$ reach-out temperature	124	(48)
$t_0$ organization onset (week)	103	(1.1)
$a_0, b_0$ shared belief prior	1.8, 129	(0.34, 23)
$\zeta$ ad-ban pseudo-successes	307	$\geq 180^\dagger$
$\chi$ rent-tier belief penalty ( $e^{-\chi} \approx 0.12$ )	2.08	(0.18)
$g$ Aug. campaign strength	0.57	(0.17)
$\lambda_{\text{cut}}$ symmetric down-cut rate	0.088	(0.006)
$m$ post-Oficio enforcement skepticism	14.8	(4.7)
<b>Panel B. Estimated model versus record</b>		
Moment	Model	Data
First coordinations (Tier 0 $\rightarrow$ 1)	214	220
Rent steps (Tier 1 $\rightarrow$ 2)	141	143
Onset / peak week	103/107	102/107
Wave spread (SD, weeks)	4.5	4.6
1 $\rightarrow$ 2 lag behind 0 $\rightarrow$ 1 (weeks)	8.1	8.0
<i>Coordination successes, by period</i>		
war / cartel	0/289	2/284
post-early / post-late	36/34	28/49
<i>Failed lifts, by period</i>		
war / cartel	110/0	123/4
post-early / post-late	46/17	46/21
Below-cost war markup	-0.09	-0.10
August war intensification ( $1+g$ )	1.57	1.68
<b>corr(log lab size, coord. week)</b>	<b>-0.38</b>	<b>-0.29</b>
Lab-wave SD (weeks)	5.8	4.9
<i>Leadership (externally disciplined calibration), CV/FA/SB %</i>		
Cartel wave (loss state, $\psi$ )	24/7/69	5/21/74
Post-early failures (loss state, $\psi$ )	6/3/91	9/9/83
Post-late rent (profit state, share)	68/20/12	65/12/22

**Note:** Panel A: eleven SMM parameters (loss 16.6,  $B = 240$ ).  $\mu$  and the tier prices  $p^x$  are in thousands of CLP ( $\hat{\mu} = 9.26 \approx 9,300$  CLP, the basket profit per drug customer). Calibrated, not estimated:  $\psi = (0.60, 0.80, 0)$  for (CV, FA, SB), shares (0.45, 0.31, 0.24),  $\tau_{\text{lead}} = 3.0$ ,  $\tau_{\text{rent}} = 0.12$ ,  $w = 26$  weeks,  $\delta = 0.95$ ,  $\tau = 0.13$ . Standard errors are the standard deviation of a 300-replication drug-resample bootstrap, each re-estimate warm-started at the headline estimate.  $^\dagger$ The bootstrap distribution of  $\zeta$  is right-skewed with a sharp floor (median 315, interquartile range [296, 370], no draw below 179) and a long upper tail, so a symmetric standard error overstates the informative uncertainty; I report the lower bound  $\zeta \geq 180$  (Online Appendix E.1). Panel B: each drug uses its own prices, cost,  $\alpha_j$ , and lab.

**The post-investigation rent.** The cartel does not break: the 220 coordinated drugs hold


their price through the FNE’s 31 March Oficio 419 (Figure 7a), and the learned enforcement belief of \$5 reproduces the contested recovery: failed lifts rebound (46 vs. near zero in the cartel) and the rent climbs late to 141 of 143 (Table 8, Panel B, skepticism  $m = 14.8$ ). The fit captures the direction but not the full sharpness: a smoothly learned belief reproduces the direction of the late surge yet undershoots its sharpness (post-late successes 34 against 49), because a smooth posterior cannot fully match a sharp freeze-then-surge. Because the held rent is a profit state, market share orders it: the dominant chain, Cruz Verde, leads the post-investigation rent, the SB→CV handoff.

Figure 5: Coordination selected by the deviation gain



**Note:** Each drug’s one-shot gain to undercutting a coordinated price,  $G_j$  (7) evaluated on the estimated demand (horizontal), against the forgone discounted escape value  $\delta \Delta V_j^1$  (9) (vertical); red = pre-ban (ads on), green = post-ban (ads banned), diamonds = regime medians, axes clipped to  $G_j \in [-3, 3]$  and  $\delta \Delta V^1 \leq 35$  (22 high-stake drugs lie above).

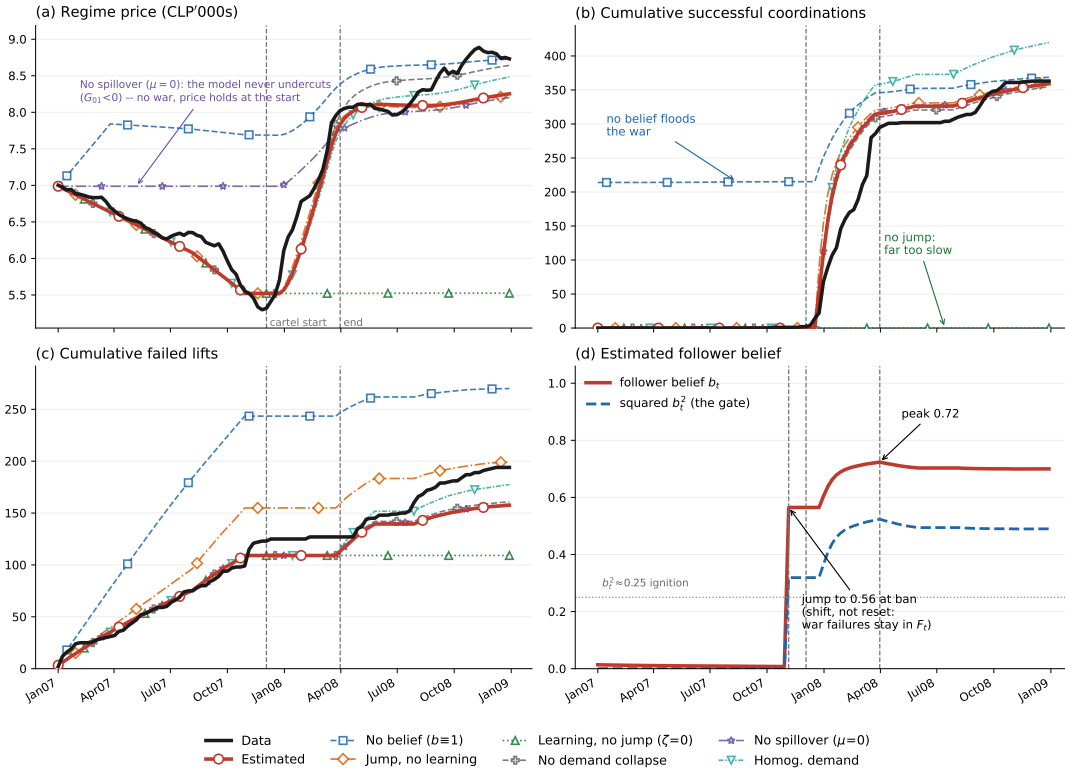
Figure 6: Model reproduces the coordination sequence



**Note: Panel A:** cumulative coordinations, model (dashed) vs. the coded data (solid), by lab-size tercile. The big- and mid-lab cohorts lead the small-lab cohort; flat through the pre-ban war, then the post-ban wave. **Panel B:** lab-level coordination week vs. log lab revenue; larger labs coordinate first (rank correlation  $-0.38$  model vs.  $-0.29$  data).

**The price path.** Figure 7a checks the model against the price series itself, with one caveat: the regime price levels are calibrated, not predicted. The coordinated tiers are the markup-implied prices of \$5, and the war floor is the pre-ban symmetric Bertrand–Nash price with store traffic, each chain maximizing  $(p - c + \mu) s(p; \alpha_{\text{pre}})$ . Because the below-cost war markup is itself a matched moment, the fitted floor reproduces the observed war price near-mechanically. Panel (a) traces the symmetric regime path against the quantity-weighted data price: the January-2007 decline to the war floor while ads are permitted, then the recovery into the coordinated tiers once the November ban flips the incentive (the regime price holds at the floor through the war, recovers from week  $\sim 102$ , and climbs to about 9 by late 2008). This is a consistency check on the whole trajectory, joined by the estimated timing, not an independent test of the level. Panels (b) and (c) make it discriminating. The cumulative wave and the failed lifts separate the estimated mechanism from the re-estimated alternatives: the belief and spillover knockouts each break a panel (the war over-produces, the wave is far too slow, or coordination collapses), whereas a misspecified homogeneous-demand spec tracks the aggregate counts but mis-shapes the wave (§5).

Figure 7: Structural specifications versus the data



**Note:** four panels, January 2007–December 2008. The data (solid black) are set against the estimated model and six knockouts of it, each switching off one ingredient: no belief ( $b \equiv 1$ ); jump, no learning; learning, no jump ( $\zeta = 0$ ); no demand collapse ( $\alpha \equiv \alpha_0$ ); no spillover ( $\mu = 0$ ); and a misspecified homogeneous demand. (a) regime price, quantity-weighted. (b) cumulative successful coordinations ( $0 \rightarrow 1$  wave and  $1 \rightarrow 2$  rent). (c) cumulative failed lifts. (d) the estimated belief  $b_t$  and its square  $b_t^2$  (the coordination gate). Each belief or spillover knockout breaks one panel; the homogeneous spec tracks the counts but mis-shapes the wave (§5, Table 9).

From one set of primitives the estimated model reproduces the main sequence: the below-cost war, the post-ban cartel wave, the post-Oficio pause, and the late rent coordination. While ads are permitted,  $G_j^{01} > 0$ , so undercutting for store traffic pays: chains cut below cost and coordinated lifts revert. The war is broad from January 2007 and intensifies with Cruz Verde’s August “Desafío” campaign (16); once the ban flips  $G_j^{01} < 0$ , undercutting ends and coordination becomes self-enforcing. The war-series and war-floor moments match the pre-ban level, August intensification, and post-ban collapse (Online

Table 9: Structural specifications versus the data: per-moment fit

Moment	Data	Estimated	No belief	Jump, no learning	Learning, no jump	No demand collapse	No spillover ( $\mu=0$ )
First coordinations (0→1)	220	214	216	214	0	191	220
Rent steps (1→2)	143	141	153	142	0	160	145
Peak week	107	107	52	108	—	107	107
Wave spread (SD, weeks)	4.6	4.5	8.3	4.8	—	5.7	4.2
<i>Coordination successes</i>							
war	2	0	<b>220</b>	0	0	0	0
cartel	284	289	137	282	0	274	294
Failed lifts, war	123	110	385	130	162	132	102
War deviations (below-cost)	516	512	<b>22</b>	536	386	509	<b>155</b>
corr(log lab size, week)	-0.29	-0.38	-0.14	-0.36	—	-0.36	-0.49
SB leads the cartel wave	0.74	0.69	0.12	0.68	—	0.64	0.70
<b>SMM criterion (loss)</b>	—	<b>16.6</b>	<b>4154</b>	<b>21</b>	<b>2323</b>	<b>28</b>	<b>612</b>

**Note:** per-moment companion to Figure 7. Each column is re-estimated under the identical SMM criterion, so every restricted spec gets its own best-fit parameters and the loss row is a fair comparison. “No demand collapse” shifts beliefs at the ban but holds price sensitivity at its pre-ban level (so  $G^{01}$  never flips); “No spillover” ( $\mu=0$ ) removes the store-traffic floor. Figure 7 additionally plots a misspecified homogeneous-demand spec (Online Appendix Table 7).

Appendix Table 4); the cumulative coordinated stock stays near zero until the ban, then rises along the wave. The estimated campaign strength  $g = 0.57$  reproduces the August intensification ( $1.57\times$  vs. the data’s  $1.68\times$ ).

The value functions sharpen the interpretation. The discounted escape value  $\delta\Delta V_j^1$  exceeds the one-shot undercut gain by more than an order of magnitude at the median drug (Figure 5), so coordination is dynamically sustainable in both regimes. Pre-ban, however, the coordinated price is not a static Nash: every chain gains by undercutting it, and the failed attempts show the chains did not coordinate on a repeated-game punishment account. The ban flips the undercut incentive to  $G < 0$  for the inelastic majority and the common-knowledge belief jump shifts expectations onto the coordinated equilibrium (the 168/52 split of §5). For the 168 static-Nash drugs, post-ban coordination is a one-shot best response throughout the rolling horizon, so there is no terminal unraveling. The remaining failed matches come from Type-I EV shocks and from the 52 non-static-Nash drugs where undercutting still pays for some chain.

## Identification and robustness

The structural exercise makes one point: the advertising ban turns the loss-leader war into self-enforcing coordination through two estimated changes together. The collapse in price sensitivity makes undercutting unprofitable for the inelastic majority; the focal belief shift releases the wave. A set of re-estimations shows the result does not rest on a tuned

parameter. Each check re-estimates the model with one ingredient switched off, plotted against the data in Figure 7. Table 9 gives the per-moment fit. Each column is re-estimated under the identical SMM criterion, so its loss row is a direct comparison. Online Appendix Table 8 adds the value-function checks.

**Robustness checks.** Belief is essential and dated to the binding November ban: a no-belief model produces far too many pre-ban coordinations (220 against the data's 2, loss 4154). Without the focal jump the wave never forms inside the sample: a learning-only belief ( $\zeta = 0$ ) coordinates essentially no drugs and fits orders of magnitude worse (2323 against the headline 16.6). A jump-only belief that drops the event-tally, re-estimated, fits essentially as well (21 against 16.6). The static flip needs the spillover: with  $\mu = 0$  the undercut gain collapses to its drug-margin term, so no chain prices below cost and the war barely forms (war deviations 155 against the data's 516; loss 612). The flip is robust. The threshold store-traffic value above which undercutting still pays rises about 3.5-fold when  $\alpha$  collapses. Before the ban it sits below the estimated  $\hat{\mu}$ , so undercutting paid and the war held; after the ban it rises above  $\hat{\mu}$ , so coordination becomes the static best response for the inelastic majority. This 3.5-fold jump stays in [3.3, 3.8] under  $\pm 10$ –30% perturbations of wholesale cost and the nesting  $\sigma$  (Online Appendix Table 10). The model is a forward-simulated Markov transition model, not a fully solved MPE; with the learned enforcement belief it reproduces the late rent (141 of 143 rent steps) and the post-investigation success/failure crossover.

**What the late rent does not capture.** The model matches the rent's total, Cruz Verde leadership, and the incentive that drives it, but not the synchronized September–October 2008 cluster: 45 of the 143 rent steps (31%) arrive together after the cartel window, during the FNE investigation. With drug-by-drug arrivals at a common rate, the model fills the same drugs gradually rather than in a burst. I report the rent's magnitude and leader, and flag this timing gap rather than let the cumulative-fit panel obscure it.

**The dynamic value function.** Three checks probe the forward-looking value (9), each re-simulated at the estimate (Online Appendix G). The model is genuinely dynamic. A myopic firm that ignores the continuation value ( $\delta \rightarrow 0$ ) coordinates only part of the static-flip set and underbuilds the rent, a loss of 555, because escaping the war and deepening the rent each trade a current cost for a future payoff. The regulatory change was also unanticipated. A perfect-foresight firm that reads the September CONAR precursor as the belief-shift date starts the wave two months early, in September (onset week 88 against the data's 102; loss 2550). The data thus place the operative shift at the binding, unforeseen November injunction, not the foreseeable precursor. Finally the continuation value must be transition-aware. Valuing each tier as an absorbing perpetuity ( $V^x = \Pi^x + \delta V^x$ ) treats a coordinated drug as if it never reverts and a rent as if it never fails. That overstates the tier values and mistimes the rent, a loss of 23, so the forward simulation over the model's own transition probabilities is what reproduces the post-investigation dynamics.

**Leadership and the loss cushion.** The chain asymmetry is the loss cushion  $\psi_i$ , calibrated from the business mix and held fixed; the leader shares are then matched by two dispersion scales. Setting  $\psi$  symmetric collapses Salcobrand's lead to chance while leaving the coordination count untouched (Online Appendix G), the signature of a leadership channel that is decoupled from the incentive constraint. Putting the asymmetry back into a chain-specific store-traffic spread  $\mu_i$  instead faces a structural bind. Store traffic enters both the escape value that selects the leader and the incentive constraint that holds the followers. If the in-sample dispersion in  $\mu_i$  is small, it cannot generate the large leadership asymmetry; if it is large enough to do so, it also makes some firms want to undercut and the coordination count breaks. The spread therefore cannot separate who leads from who follows. The loss cushion separates these margins: it moves leadership through the exposure-scaled choice value while leaving enough slack that the coordination count barely moves. The share magnitudes are calibrated to the data (Salcobrand leads 74% of cartel-wave increases, the model 69%), but the ordering comes from business mix and ownership rather than from prices: FASA's Mexico/Peru operations, Cruz Verde's Socofar distribution structure, and Salcobrand's domestic pharmacy focus and 2007 ownership change. Fixing the two dispersions at a common value still leaves Salcobrand leading the wave (69–80% across the range, Online Appendix G.3). Freeing  $\psi_{SB}$  and re-estimating returns it to zero, supporting the calibrated most-exposed corner; the August 2007 acquisition and management change at Salcobrand, which inherited the intensifying below-cost war and led the December escape, is its institutional reading.

**If prices are deflated for inflation.** Deflating prices to the war base for the 2007–2008 inflation leaves the war below cost and the cartel, the rent, and the leadership intact (Online Appendix G.1).

**If the demand estimate is wrong.** Because the whole flip is read off the estimated demand, I re-estimate the structural model under alternative demand specifications (Online Appendix Table 7): the per-drug heterogeneous  $\hat{\alpha}_j$  specifications reproduce the leadership and the count, while a homogeneous demand at the median  $\hat{\alpha}$  fits worse because, without the per-drug heterogeneity, the wave is somewhat too concentrated and the rent is too deep and too slow.

## Welfare of the equilibrium transition

What did the transition cost consumers? Because the war was below-cost and the cartel a coordinated increase, the answer depends on the counterfactual. I evaluate the settled coordinated prices of January–April 2008 (fourteen weeks, the back of the nineteen-week wave) against two benchmarks: the below-cost war prices of mid-2007 (what consumers faced when the cartel formed) and the pre-war 2006 baseline, computing for each drug the consumer transfer and, from the nested logit, the deadweight loss from the induced

quantity distortion.<sup>25</sup>

The measure is within the pharmaceutical market: the drug overcharge and the quantity-distortion deadweight loss, both identified by the estimated drug demand. The coordinated and benchmark price levels are read directly from the data (the same median tier prices the structural model is calibrated to), so the transfer is the realized overcharge, not a price path the model predicts. I do not extend it to the non-pharmaceutical basket, which the case record cannot price; the coordination raised drug prices, so this counts the pharmaceutical overcharge alone.

The coordinated path extracted supra-competitive rents on many products (Table 10). Within the cartel window 37% of the coordinating drugs reached Tier 2 rent extraction, deepening to 143 drugs (65% of all 220) by year-end as it ratcheted post-Oficio. These rents were priced at a median 1.48× wholesale cost and 1.11× the 2006 baseline (11% above it), the margins the conduct was fined for [Fiscalía Nacional Económica (FNE), 2008, Tribunal de Defensa de la Libre Competencia (TDLC), 2012].

Table 10: Welfare of the coordinated price path

<b>Panel A. Coordinated prices above margin</b>	
Drugs reaching Tier 2 within the window	37%
Median Tier-2 price / wholesale cost	1.48×
Coordination events priced > 1.4× cost	44%
Median Tier-2 price / 2006 baseline (% above)	1.11× (11%)
<b>Panel B. Consumer harm over the window</b>	
Transfer vs. below-cost war (% of sales)	+CLP 3.4bn (+11%)
of which from the Tier-2 rent drugs	~CLP 2.7bn (80%)
Gross rent vs. 2006 baseline (above-historical drugs)	+CLP 1.1bn
Deadweight loss (quantity distortion)	<CLP 0.03bn (< 0.1%)

**Note:** Over the settled fourteen-week window (Jan–Apr 2008, three chains; 222 drugs with complete baseline, war, and cartel price windows; window sales ≈CLP 30.7bn). Panel A: how far coordinated prices stood above margin. Panel B: consumer harm, the overcharge on the transacted quantity (transfer) plus the deadweight loss from the quantity response in the estimated nested logit. The transfer and the deadweight loss are deliberately distinct objects (the first the overcharge incidence on the transacted quantity, the second the efficiency loss from the induced quantity distortion), reported separately, not summed.

<sup>25</sup>Per drug, with the estimated nested logit ( $\sigma = 0.393$ ; post-ban  $\alpha = 0.029$ , the same demand the structural model uses). The consumer-surplus change between the coordinated price  $p^c$  and benchmark  $p^b$  is  $M [\text{CS}(p^b) - \text{CS}(p^c)]$ , with  $\text{CS}(p) = \alpha^{-1} \log(1 + [N e^{(\phi - \alpha p)/(1 - \sigma)}]^{1 - \sigma})$ , decomposed into the transfer  $(p^c - p^b) \times$  quantity and the residual deadweight loss;  $M$  is backed out from the cartel-window quantity and the demand-implied inside share at the cartel price; here  $N = 3$  counts the inside chains, not the market size  $N_j$  of (4), which enters only through  $M$ .

Relative to the below-cost war prices consumers had been paying, the coordinated path cost them about 11% (CLP 3.4bn over the fourteen-week window), four-fifths of it from the Tier-2 rents. The FNE's gross-revenue figure over the full conduct and the TDLC fine are on a different time frame or basis (the fine is a deterrence penalty, not a measure of consumer harm), so a fourteen-week transfer is not directly comparable to them [Fiscalía Nacional Económica (FNE), 2008, Tribunal de Defensa de la Libre Competencia (TDLC), 2012]. The deadweight loss is small ( $< \text{CLP } 0.03\text{bn}$ , under 0.1% of sales): post-ban demand is inelastic, so the increase moved little quantity and the harm is overwhelmingly a transfer. The extraction was selective: the cartel held the high-volume, salient staples near competitive levels (volume-overcharge correlation  $-0.14$ ) while taking rents on the less-salient products.

## 6 Conclusion

I have analysed the equilibrium transition from loss-leader competition to coordinated pricing among the three large Chilean retail pharmacy chains in 2007–2008. The analysis yields three results.

First, the transition needs two changes at once, and neither alone is sufficient. The war was a low-belief trap: each chain expected a price rise to be undercut, so none held one. The advertising ban broke the trap in two ways. As a public, dated event it made the chains expect a coordinated rise to hold this time. And by removing the comparative-price advertising that told shoppers who was cheapest, it cut the estimated price coefficient from 0.103 to 0.029, so undercutting won too little traffic to repay the margin loss. Once undercutting stopped paying, the coordinated price held on its own for those drugs, with no patience and no off-path discipline, and the rolling-horizon terminal condition did not unravel it. The belief jump and the demand change are each necessary; only together do they reproduce the path (§5). The record supports the interpretation but does not estimate the belief: it documents the unilateral-demand-diversion problem and the laboratory-mediated verification of the agreement's success [Fiscalía Nacional Económica (FNE), 2008, Tribunal de Defensa de la Libre Competencia (TDLC), 2012].

Second, the cartel's leadership and its slow rollout are disciplined by external primitives, not by a free chain-specific spillover. The case file records who led, and that the laboratories organized the coordination, but it does not say why. The model does. Salcobrand leads because it has the least non-drug business to cushion the below-cost war, so it has the most to gain from ending it. The larger laboratories move first because they carry the biggest prize, so the model's reach-out engages them earliest. Both come from primitives I calibrate outside the model, the chains' loss cushions and the labs' stakes, and the model matches the leadership, the larger-labs-first sequence, the war, the post-ban

wave and its February pause, and the split between margin restoration and rent. The post-ban failures are mostly stray unmatched lifts and reverted rent pushes, rather than defections from coordinated prices.

Third, the harm is a transfer, not an efficiency loss. The coordinated path raised prices well above cost on many products; 37% reached rent prices near 1.5× wholesale cost, costing consumers about CLP 3.4bn relative to the below-cost war prices over the cartel window, four-fifths of it on the rent products. Because post-ban demand is inelastic, the deadweight loss is tiny. The cartel mainly moved money from consumers to firms, and it did so where consumers paid least attention.

The model's economic content rests on objects estimated or documented outside it: the demand system, the wholesale costs and tier prices, and the laboratory structure and leadership from the legal record. The store-traffic value is bounded by the chains' own below-cost behaviour, and a handful of organizational-timing parameters are calibrated by simulated method of moments. The gradual, four-month rollout, which the prosecutor's complaint describes as expanding as the agreement's success was verified,<sup>26</sup> is reproduced by the endogenous laboratory reach-out rather than by a separate contagion term. More chains find it worth organizing a laboratory as undercutting stops paying and beliefs shift, drug by drug. What gradual belief-learning adds beyond this is a secondary refinement (§5). Disentangling it from antitrust-detection risk and operational complexity is a useful direction for future work.

## References

- Victor Aguirregabiria. The dynamics of markups and inventories in retailing firms. *The review of economic studies*, 66(2):275–308, 1999.
- Victor Aguirregabiria and Jihye Jeon. Firms' beliefs and learning: Models, identification, and empirical evidence. *Review of Industrial Organization*, 56(2):203–235, 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.

---

<sup>26</sup>Original: “a medida que se verificaba el éxito del acuerdo” [Fiscalía Nacional Económica (FNE), 2008].

- Patrick Bajari, C Lanier Benkard, and Jonathan Levin. Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370, 2007.
- Lee Benham. The effect of advertising on the price of eyeglasses. *Journal of Law and Economics*, 15(2):337–352, 1972.
- Steven T Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262, 1994.
- Jonas Björnerstedt and Frank Verboven. Does merger simulation work? Evidence from the Swedish analgesics market. *American Economic Journal: Applied Economics*, 8(3):125–164, 2016.
- David P. Byrne and Nicolas de Roos. Learning to coordinate: A study in retail gasoline. *American Economic Review*, 109(2):591–619, 2019.
- John F. Cady. An estimate of the price effects of restrictions on drug price advertising. *Economic Inquiry*, 14(4):493–510, 1976.
- Juan Esteban Carranza, Robert Clark, and Jean-François Houde. Price controls and market structure: Evidence from gasoline retail markets. *The Journal of Industrial Economics*, 63(1):152–198, 2015.
- 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.
- Steve Cicala. When does regulation distort costs? lessons from fuel procurement in us electricity generation. *American Economic Review*, 105(1):411–444, 2015.
- Robert Clark and Jean-François Houde. Collusion with asymmetric retailers: Evidence from a gasoline price-fixing case. *American Economic Journal: Microeconomics*, 5(3):97–123, 2013.
- Robert Clark and Jean-François Houde. The effect of explicit communication on pricing: Evidence from the collapse of a gasoline cartel. *The Journal of Industrial Economics*, 62(2): 191–228, 2014.
- Comisión Resolutiva H. de Defensa de la Libre Competencia. Resolución n° 432/1995: Caso cruz verde – salcobrand – farmacias ahumada. Reproduced by CIPER Chile, 1995. <https://www.ciperchile.cl/2009/03/31/documento-el-fallo-que-sanciona-a-las-farmacias-por-colusion-en-1995/>.

- Christopher T. Conlon and Jeff Gortmaker. Best practices for differentiated products demand estimation with PyBLP. *The RAND Journal of Economics*, 51(4):1108–1161, 2020.
- Patrick DeGraba. Volume discounts, loss leaders and competition for more profitable customers. Working Paper 260, Federal Trade Commission, Bureau of Economics, 2003.
- Diario Financiero. Empresas Juan Yarur (familia Yarur–Abuhadba) adquiere Salcobrand por US\$130 millones, agosto 2007 — reestructuración hasta 2008. Press archive. Cross-referenced in Tribunal de Defensa de la Libre Competencia [2012] (§ “Estructura accionaria de Salcobrand”) and CMF emisor record Empresas Juan Yarur SpA, RUT 91717000, 2007.
- Andrés Díaz and Alexander Galetovic. El mercado de las farmacias en Chile: ¿competencia por localización? *Economía y Política*, 2(1):33–73, 2015. <https://dialnet.unirioja.es/descarga/articulo/6266396.pdf>.
- 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.
- E. Woodrow Eckard. Competition and the cigarette TV advertising ban. *Economic Inquiry*, 29(1):119–133, 1991.
- Sara Fisher Ellison, Iain Cockburn, Zvi Griliches, and Jerry Hausman. Characteristics of demand for pharmaceutical products: An examination of four cephalosporins. *RAND Journal of Economics*, 28(3):426–446, 1997.
- Chaim Fershtman and Ariel Pakes. Dynamic games with asymmetric information: A framework for empirical work. *The Quarterly Journal of Economics*, 127(4):1611–1661, 2012.
- Fiscalía Nacional Económica. Estudio de mercado sobre medicamentos (EM03-2018): Informe Final. Technical report, FNE Chile, 2020. <https://www.fne.gob.cl/wp-content/uploads/2020/01/Informe-Final.pdf>.
- Fiscalía Nacional Económica. Investigación del mercado de productos farmacéuticos, Rol 1723-10 — estructura del mercado, integración vertical Socofar/Cruz Verde, canales de distribución. Informe publicable. URL: [http://www.fne.cl/transparencia/libro/archivos/archivo/informe\\_1723-10.pdf](http://www.fne.cl/transparencia/libro/archivos/archivo/informe_1723-10.pdf), 2010.
- Fiscalía Nacional Económica (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](https://www.fne.gob.cl/wp-content/uploads/2011/03/requ_0009_2008.pdf). Legal complaint, Tribunal de Defensa de la Libre Competencia.

- Jorge Florez-Acosta and Daniel Herrera-Araujo. Multiproduct retailing and consumer shopping behavior: The role of shopping costs. *International Journal of Industrial Organization*, 68:102560, 2020.
- 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.
- Avi Goldfarb and Mo Xiao. Who thinks about the competition? managerial ability and strategic entry in us local telephone markets. *American Economic Review*, 101(7):3130–61, 2011.
- Dana P. Goldman, Geoffrey F. Joyce, José J. Escarce, Jennifer E. Pace, Matthew D. Solomon, Marianne Laouri, Pamela B. Landsman, and Steven M. Teutsch. Pharmacy benefits and the use of drugs by the chronically ill. *JAMA*, 291(19):2344–2350, 2004.
- Henry G Grabowski and John M Vernon. Brand loyalty, entry, and price competition in pharmaceuticals after the 1984 drug act. *The journal of law and economics*, 35(2):331–350, 1992.
- Edward J. Green and Robert H. Porter. Noncooperative collusion under imperfect price information. *Econometrica*, 52(1):87–100, 1984.
- Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054, 1982.
- Jr. Harrington, Joseph E. Lectures on Collusive Practices. CRESSE lecture, University of Pennsylvania, 2018.
- Jerry A. Hausman. Valuation of new goods under perfect and imperfect competition. In Timothy F. Bresnahan and Robert J. Gordon, editors, *The Economics of New Goods*, pages 207–248. University of Chicago Press, 1996.
- Yufeng Huang, Paul B Ellickson, and Mitchell J Lovett. Learning to set prices. *Journal of Marketing Research*, 59(2):411–434, 2022.
- Humphreys Clasificadora de Riesgo. Informe de clasificación (Anual): Farmacias Ahumada S.A., octubre 2007 — mix ebitda México/chile/Perú, mercado Chile, presión competitiva supermercados y perfumerías. Rating report. URL: <https://humphreys.cl/wp-content/uploads/2018/08/fasaoct07anual.pdf>, 2007.
- Mitsuru Igami. Estimating the innovator’s dilemma: Structural analysis of creative destruction in the hard disk drive industry, 1981–1998. *Journal of Political Economy*, 125(3): 798–847, 2017.

- 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.
- Kazuko Kano. Menu costs and dynamic duopoly. *International Journal of Industrial Organization*, 31(1):102–118, 2013.
- Christopher R Knittel and Victor Stango. Price ceilings as focal points for tacit collusion: Evidence from credit cards. *American Economic Review*, 93(5):1703–1729, 2003.
- Rajiv Lal and Carmen Matutes. Retail pricing and advertising strategies. *Journal of Business*, 67(3):345–370, 1994.
- Matthew S Lewis. Odd prices at retail gasoline stations: focal point pricing and tacit collusion. *Journal of Economics & Management Strategy*, 24(3):664–685, 2015.
- Daniel McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026, 1989.
- Nathan H Miller and Matthew C Weinberg. Understanding the price effects of the miller-coors joint venture. *Econometrica*, 85(6):1763–1791, 2017.
- 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.
- Ariel Pakes and David Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057, 1989.
- Ram C Rao and Niladri Syam. Equilibrium price communication and unadvertised specials by competing supermarkets. *Marketing Science*, 20(1):61–81, 2001.
- Stephen P Ryan. The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061, 2012.
- Michael Sinkinson and Amanda Starc. Ask your doctor? direct-to-consumer advertising of pharmaceuticals. *The Review of Economic Studies*, 86(2):836–881, 2019.
- James H. Stock and Motohiro Yogo. Testing for weak instruments in linear IV regression. In Donald W. K. Andrews and James H. Stock, editors, *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, pages 80–108. Cambridge University Press, 2005.

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.

Tribunal de Defensa de la Libre Competencia. Sentencia n° 119/2012, rol c 184-2008: Requerimiento de la Fiscalía Nacional Económica contra Farmacias Ahumada S.A., Farmacias Cruz Verde S.A. y Salcobrand S.A. (333pp; margen sugerido 20-25%; ganancias del cártel 2.8%/4.0%/3.1% ventas; 222 medicamentos  $\approx$  16% ingresos). Sentencia de 31 de enero de 2012. URL: [https://www.fne.gob.cl/wp-content/uploads/2012/01/Sentencia\\_119\\_2012.pdf](https://www.fne.gob.cl/wp-content/uploads/2012/01/Sentencia_119_2012.pdf), 2012.

Tribunal de Defensa de la Libre Competencia (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](https://www.fne.gob.cl/wp-content/uploads/2012/01/Sentencia_119_2012.pdf).