Online Appendix to "Equilibrium Transition and Cartel Formation: A Structural Analysis of Chile's Pharmacy Cartel"

Yu (Jasmine) Hao

September 15, 2026

This Online Appendix accompanies the main text and contains the case-record evidence. Chain abbreviations are Cruz Verde (CV), Farmacias Ahumada (FASA; FA in equations and tables), and Salcobrand (SB). Institutional abbreviations are Chile's National Economic Prosecutor's Office (*Fiscalía Nacional Económica*, FNE), the Competition Tribunal (*Tribunal de Defensa de la Libre Competencia*, TDLC), and the advertising self-regulation council (*Consejo de Autorregulación y Ética Publicitaria*, CONAR). Monetary amounts are in Chilean pesos (CLP).

A Industry background and case records

(a) Sources and evidentiary conventions. This appendix reports the institutional facts and case-record evidence used in the article. It draws on the following sources: the FNE's *Requerimiento* of 9 December 2008; the TDLC's *Sentencia* No. 119/2012 of 31 January 2012; the court-appointed expert report of 26 October 2010; the Supreme Court's judgment of 7 September 2012; a contemporary account of the advertising dispute by Chile's consumer-protection agency (SERNAC); and Cruz Verde's court filings, which reproduce the civil injunction and testimony about its effects (FNE, 2008; TDLC, 2012; Núñez et al., 2010; Corte Suprema de Chile, 2012; Servicio Nacional del Consumidor, 2007; Farmacias Cruz Verde S.A., 2009, 2012). Item-specific academic, press, and article-data sources appear where used. The proceedings were assembled to decide liability. Their account of the pricing mechanism is incidental to that legal purpose. I use the cited passages to support the model's institutional assumptions.

I distinguish compliance monitoring, batch success, and procedure verification. Compliance monitoring asks whether a proposed price has been implemented. A batch is a list of

medicines proposed for coordinated price increases. I classify a batch as successful when all three chains reach each medicine's proposed price. The first complete laboratory-mediated implementation supplies a timing benchmark for procedure verification. In the structural model, verification follows a weekly information transition . The observed implementation does not directly reveal the followers' information state.

Each Spanish excerpt follows an English account of its meaning and identifies its source. Quotations from the judgment may reproduce either the tribunal's own words or testimony quoted by the tribunal. I used AI assistance to prepare the English translations . The original Spanish controls. For the FNE complaint, I cite the numbered paragraph and page. The TDLC judgment contains two parts: *Vistos*, which records the parties' submissions and the procedural history, and numbered *Consideraciones*, which state the tribunal's findings. References to *Vistos* give the section number; references to *Consideraciones* use "consid." followed by the numbered finding.

Industry background

(b) Three chains setting one national price schedule. In 2007, the three chains' joint share of national retail medicine sales was 92% and was uncontested in the proceeding (TDLC, 2012, consid. 54). The FNE puts the 2007 figure at 92%, citing IMS Health Chile (FNE, 2008, para. 37, p. 17) (TDLC, 2012, *Vistos* 1.20). Cruz Verde was vertically integrated with the distributor Socofar, which from October 2007 also supplied generics to Salcobrand (FNE, 2008, paras. 8 and 36, pp. 4 and 16–17).

Each chain centrally administered a national price schedule, with exceptions for remote regions (Fiscalía Nacional Económica, 2011, paras. 513–515, pp. 179–180). Salcobrand's pricing rules explicitly referred to its rivals. Its commercial manager described the standing policy for non-salient products as being "equal to Cruz Verde," and proposed moving to 3% above Cruz Verde on all products while staying 1% below FASA on the salient ones (TDLC, 2012, consid. 95).

(c) The FNE's account of prescription constraints. The FNE argued that prescriptions limited consumers' ability to substitute other medicines. It described consumers as captive through the prescription and the affected medicines as difficult to substitute. The judgment records this submission in *Vistos* 1.25, and the respondents disputed it.

> *"Consumidores cautivos a través de la receta médica, característica que los hace difícilmente sustituibles."*
>
> (FNE, 2008, para. 74, p. 27)

(d) The PVPS was a common, manufacturer-supplied benchmark. Manufacturers sent all three chains price lists carrying a suggested retail price, the PVPS. According to the FNE, the suggested price left the pharmacy chains a margin of 20–25% and reflected the product's position within its therapeutic category.[1]

> *"Dichos PVPS se determinan considerando un margen de entre un* 20% *y un* 25% *para las Cadenas de Farmacias, y el posicionamiento del producto dentro de su categoría terapéutica."*
>
> (FNE, 2008, para. 13, p. 5)

Paragraph 96 repeats this measure and notes that the practice of circulating suggested retail prices had been held an unlawful commercial practice by the Comisión Preventiva Central in 1983 (Dictamen No. 387 of 20 July 1983) (FNE, 2008, para. 96, p. 35).

(e) Temporary discounts became persistent cuts at the end of 2006**.** Through most of 2006 discounts reverted. Prices fell during promotions and returned to their previous level. Figure 1 shows the pattern and its break. Salcobrand's price record breaks on 22 October 2006 and Cruz Verde's on 23 November 2006 . After those dates discounts became more persistent and prices fell through 2007. The price data identify when the change occurred, but the case record does not explain why temporary discounts became persistent cuts.

[1]The cited FNE passages do not specify whether this margin is measured relative to cost or the retail price.

Figure 1: Scheduled markdowns, reversals, and persistent low prices

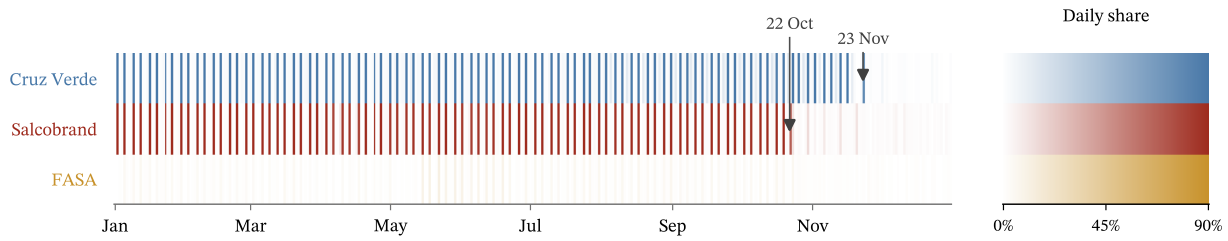

(a) Daily coverage of price cuts of 15% or more

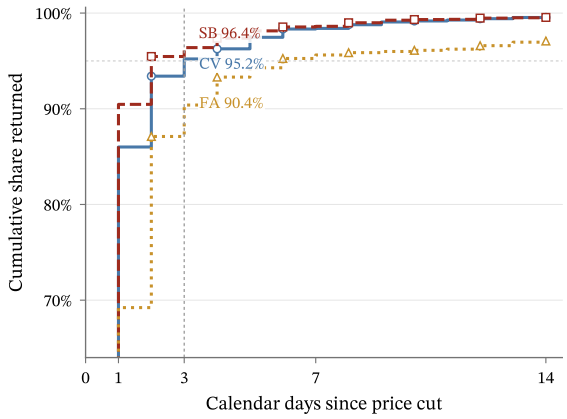

(b) Share returning to the pre-cut price

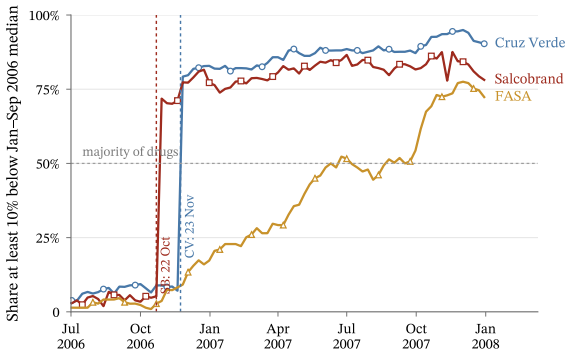

(c) Transition to persistent low prices

Note: Panel (a) plots, for each chain and day, the share of medicines priced at least 15% below the preceding calendar day's price, excluding drops to less than one-third of that price. The denominator includes medicines with valid prices on both days . Color intensity indicates the share. Panel (b) uses markdown spells beginning in January–September 2006 and plots the Kaplan–Meier estimate of the cumulative probability of returning to at least 95% of the pre-cut price, accounting for right-censored spells. Panel (c) reports the weekly share of medicines whose weekly median price is at least 10% below their baseline, defined as the median of weekly median prices in January–September 2006. Its denominator contains medicines with a valid weekly price and baseline and can vary across chain-weeks. A weekly price requires at least four daily observations. The horizontal line marks 50% . The vertical lines mark Salcobrand's 22 October cut and Cruz Verde's 23 November cut.

(f) Comparative advertising intensified the price war in August 2007. In August 2007 Cruz Verde launched, through television and print advertising, the Cruz Verde Challenge, whose slogan promised low prices without competition. The campaign compared Cruz Verde's prices on 685 high-turnover branded medicines with FASA's prices for the same products.

> *"Desafío Cruz Verde, Precios Bajos Sin Competencia," en que "comparó los precios de* 685 *medicamentos de marca y de alto volumen de rotación, con los precios de los mismos productos en Fasa."*
>
> (FNE, 2008, para. 3, pp. 2–3)

That the chains fought this price war, that it compressed margins for all three, and that it produced below-cost sales were uncontested in the proceeding (TDLC, 2012, consid. 59).

(g) The advertisement ban became binding on 6 **November** 2007**.** FASA challenged the campaign before CONAR, the advertising industry's self-regulator. SERNAC reported on 31 October 2007 that CONAR had upheld its sanction but that self-regulation had not stopped the campaign (Servicio Nacional del Consumidor, 2007). FASA also sued for unfair competition in Santiago's 17th Civil Court on 25 October (FNE, 2008, para. 3, p. 3, n. 2). Cruz Verde's answer to the FNE complaint records that the court granted and notified the advertisement ban on 6 November 2007. It reproduces the order suspending the campaign and other comparative advertising against FASA pending the court's decision (Farmacias Cruz Verde S.A., 2009, p. 17 and n. 5).

The ban stopped Cruz Verde from advertising its price advantage over FASA. In testimony reproduced in Cruz Verde's appeal, FASA's chief executive, Sergio Purcell, said that, in his view, it no longer made sense for Cruz Verde to continue escalating price cuts if it could not advertise them (Farmacias Cruz Verde S.A., 2012, p. 110).

> *"deja de tener sentido, a mi modo de ver para Cruz Verde seguir con la escalada de reducción de precios si no la podía publicitar".*
>
> (Farmacias Cruz Verde S.A., 2012, p. 110)

(h) Salcobrand's new management sought to restore margins. Against the price-war losses documented in (f), Salcobrand, which changed ownership and management in 2007, stated in its defence that the new management raised margins on negative-margin products even at the risk of losing market share. It described this strategy as necessary for the firm's long-run viability.

> *"Incluso a riesgo de perder participación en el mercado"; "se justificaba como una necesidad para la viabilidad de la empresa en el largo plazo."*
>
> (TDLC, 2012, *Vistos* 2.3)

Case records

(i) Who proposed the coordination. The parties and the laboratory managers described their roles differently. FASA's settlement states that some of its executives met

personally with laboratory executives in November 2007. Some laboratory executives then conveyed a proposal for all three chains to raise prices jointly on a specified group of medicines as a solution to the prevailing market conditions.

> *"En noviembre de* 2007*, algunos ejecutivos de FASA mantuvieron contactos personales con ejecutivos de algunos laboratorios." "Algunos de tales ejecutivos de laboratorios transmitieron a los ejecutivos de FASA la proposición de alzar coordinadamente los precios de las tres compañías (FASA, Salcobrand y Cruz Verde) para un grupo determinado de medicamentos, como solución a esta situación de mercado."*
>
> (TDLC, 2012, consid. 12, 63)

Recalcine's sales manager described requests from the chains to the laboratory. He said that the chains asked him to coordinate the dates of price increases so that no chain would be left less competitive. The laboratory benefited because this work kept open a channel through which it could ask a chain to lower a price that stood far above the market leader's. The chains also held him responsible for successful coordination even though, he said, he lacked authority to make the other chains comply.

> *"una u otra quede menos competitiva"; "tiene las puertas abiertas"; "De alguna forma responsabilizan en mí el éxito de la coordinación, aunque yo no tenga poder de hacerla cumplir al resto de las cadenas."*
>
> (TDLC, 2012, consid. 85)

The tribunal resolved the conflict against Salcobrand based on an internal Salcobrand email dated 19 December 2007. The commercial manager listed the actions taken to reverse rivals' price cuts. The third was to insist that laboratories coordinate price increases. The tribunal read the verb "insist" as evidence of an established practice and placed the initiative with the pharmacies.

> *"Insistir con los laboratorios la necesidad de una coordinación para el alza de sus precios." "La iniciativa de coordinación habría sido de las farmacias, y no de los laboratorios."*
>
> (TDLC, 2012, consid. 95, 97)

(j) What the laboratories transmitted. The FNE states that the chains used laboratories to coordinate and monitor the agreement. Laboratories conveyed the size and date of a price increase several days in advance, sometimes identified which chain should move first, and later reported failures of coordination.

> *"utilizaron a los laboratorios para coordinar y monitorear el acuerdo, al punto que los laboratorios llegaron a servir para comunicar, con varios días de anticipación, la entidad del alza del precio, su fecha de ocurrencia y, en casos puntuales, cuál de ellas debería reflejar primero el alza, e incluso después para representar descoordinaciones"*
>
> (FNE, 2008, para. 16, p. 6)

One contemporaneous email and one laboratory-manager statement provide direct evidence of the same role. On 17 December 2007, the general manager of Laboratorio Medipharm wrote to Ramón Ávila at Salcobrand, listing four medicines whose final retail prices would rise. He asked Salcobrand to identify Day One so that he could inform the others and requested a message on his mobile phone. Asked to explain the email, he said that Day One was his own expression for launching a promotional campaign, that there was no Day Two, and that "the others" were his medical representatives. Mehilín Velásquez, to whom Ávila forwarded the email, did not explain the request for a mobile message (TDLC, 2012, consids. 125–126). The tribunal rejected these alternative explanations (TDLC, 2012, consid. 128). The reference to medical representatives appears only in that explanation; the people who relayed prices and dates in the testimony were laboratory general managers, sales managers, and key-account managers.

> *"[n]ecesito saber cuál será el DÍA* 1 *para informarle a los demás, avísame al celular"; "DÍA* 1*"; "una manera de hablar propia mía"; "el día* 2 *no existe"; "los demás"; "mi gente, los visitadores médicos con quienes yo hago las campañas".*
>
> (TDLC, 2012, consid. 125)

The laboratory-manager statement describes the same role from the other side. The key-account manager of Laboratorio Grünenthal stated that Ahumada asked him to contact Cruz Verde about raising the price of Belara. He did so, the increase was coordinated, and its implementation was checked and reported to the chains.

> *"Ahumada me pidió que hablara con CV para que subiera los precios de Belara, y yo hablé y se coordinaron para subirlo"; "esto fue chequeado y se comunicó a las cadenas que había subido el precio conforme a lo acordado".*
>
> (TDLC, 2012, consid. 86)

(k) The one-two-three procedure: Day One, Day Two, Day Three. FASA's deputy manager for prescription pharmaceuticals described the mechanism in detail. In November or early December 2007, Laboratorio Bayer proposed that Salcobrand raise first the prices of a list of gynaecological products supplied by the laboratory, with FASA and Cruz Verde following. On Day One, Salcobrand raised the whole list. On Day Two, FASA raised one half and Cruz Verde the other. On Day Three, each follower raised the half it had not raised the day before. The witness called the mechanism Day One, Day Two, and Day Three.

"Día Uno"; "Día Dos"; "Día Tres".

(TDLC, 2012, consid. 75)

The same 19 December email quoted in (i) contains Salcobrand's account of the design. Having listed insisting with the laboratories as the third action taken, the manager continued that Salcobrand had offered to raise prices first, on a Monday or Tuesday, leaving the other two chains three or four days to detect and adopt the increase. The tribunal noted that the words "detect" and "procedure" appear in quotation marks in the original and read both as euphemisms.

"Para ello ofrecimos ser la cadena que primero subiera los precios (los días lunes o martes) de este modo, las otras dos cadenas tendrían 3 o 4 días para 'detectar' estas alzas y luego asumirlas".

(TDLC, 2012, consid. 95, 97)

The following evidence describes the timing. The tribunal found that the pattern it called the collusive mechanism "1-2-3" had Salcobrand generally raising the relevant price list first and the other two following, half the time Cruz Verde first and half the time FASA. The sequence appeared in the market within three or four days. Some later sequences departed from this order, and FASA sometimes moved first.

"que se veía reflejado en el mercado en 3 ó 4 días".

(TDLC, 2012, consid. 128)

A FASA witness placed the starts on Monday through Wednesday. This timing let all three chains implement an increase and avoided leaving one chain alone with the higher price over the weekend. The tribunal found that 85% of the identified increases began on a Monday, Tuesday, or Wednesday (TDLC, 2012, consid. 74).

"para asegurarnos que todos iban a hacerlo y que no se iba a quedar una sola cadena de farmacias con el alza y con el fin de semana de por medio que le podría afectar en su competitividad".

(TDLC, 2012, consid. 72)

The case record reports three counts of sequential increases, each based on a different definition. The definitions change the required timing, persistence, and frequency of rival-price checks (quotation intensity). The court-appointed experts define a one-two-three increase by three conditions: all three chains raise the price; one raises first and within at most four days the other two raise, on different days; and the resulting prices hold for at least three days. On that definition they count 32 such increases in January 2006–November 2007, 162 in December 2007–April 2008, and 21 in May–December 2008, or 1.39, 32.40 and 2.63 per

month (Núñez et al., 2010, pp. 51–52, Tables IV.4–IV.5). Within the middle period the two most frequent orders are Salcobrand–FASA–Cruz Verde (85 of 162) and Salcobrand–Cruz Verde–FASA (65 of 162). These two Salcobrand-first orders account for 150 of 162 increases (92.6%; reported as 92% by the experts). Over the full 2006–2008 sample, the experts count Salcobrand first in 162 of 215 increases (reported as 75%). The judgment repeats the two middle-period order counts but pairs its following 75% statement with the middle-period 162-event denominator, combining the report's full-sample percentage with its middle-period denominator (TDLC, 2012, consid. 141).

The one-two definition requires both followers to raise prices on the same later day, within two days of the first increase, with the resulting prices holding for at least three days (TDLC, 2012, consid. 142). The same experts separately count one-two increases, in which the two followers move together on a single later day. In December 2007–April 2008, Salcobrand leads 15 of 29 and FASA 14, and across the whole sample only three increases had all three chains moving on the same day and holding (TDLC, 2012, consid. 143).

In plain language, the quotation-intensity definition drops the four-day timing and three-day persistence requirements. Finally, the tribunal tabulated the December 2007–March 2008 increases among the 222 medicines that coincide with a spike in FASA's and Cruz Verde's quotation intensity and omit the four-day and three-day conditions: 125 Salcobrand–FASA–Cruz Verde, 105 Salcobrand–Cruz Verde–FASA, 12 and 7 with FASA first, and 37 increases in other orders, 286 in all (TDLC, 2012, consid. 188, Table 262).

(l) Monitoring: prior knowledge and quotation intensity. The FNE describes two monitoring channels. The chains conducted periodic competitor price checks using their own or contracted staff. Laboratories also tracked the retail prices of their products and reported the results to the chains.

> *"a. Cotizaciones periódicas: Efectuadas por las requeridas, ya sea con personal propio o externo." "b. Monitoreo a través de los laboratorios farmacéuticos, los cuales hacen seguimiento de los precios de venta a público de sus productos ... y cuyos resultados eran comunicados a las Cadenas de Farmacias."*
>
> (FNE, 2008, para. 96, p. 35)

FASA's witness said that the chain monitored rival compliance by adding extraordinary price surveys to its regular checks and increasing the survey budget.

> *"encuestas extraordinarias—es decir, adicionales a las realizadas con regularidad—aumentando el presupuesto para ello".*
>
> (TDLC, 2012, consid. 75)

One exchange shows compliance monitoring directly. At 9:53 on 14 March 2008, Judith Carreño of Salcobrand told Laboratorio Gardenhouse that the price would change that day and asked it to verify that competitors had also changed theirs. Eighty-eight minutes later, the laboratory reported that Ahumada's new price was active in stores and that Cruz Verde's would appear that afternoon after a transmission delay. The tribunal held that this reply was incompatible with mere price following because it conveyed what another chain would do in the future.

> *"[h]oy se modificará el precio pero verifica que la competencia lo hubiese cambiado también"; "Ahumada lo cambió y está listo en locales. Cruz Verde también pero tuvo un retraso en la transmisión y va a aparecer durante la tarde"; "lo que sólo se explica en el marco de una colusión".*
>
> (TDLC, 2012, consid. 130)

The quotation data give the same result with a baseline. The tribunal deliberately excluded Salcobrand's own quotations from the comparison, on the ground that the first mover's incentives to monitor rivals should be similar with and without collusion. The informative variation therefore lies in the followers.

> *"los incentivos ... para monitorear el comportamiento de su competencia, debieran ser similares con y sin colusión".*
>
> (TDLC, 2012, consid. 173)

Cruz Verde normally quoted each medicine at FASA two to four times a week and at Salcobrand less than once a week. During the increases, it quoted at both chains for several consecutive days, beginning at Salcobrand on the day of, or the day before, Salcobrand's increase (TDLC, 2012, consid. 174). The tribunal concluded that Cruz Verde and FASA knew in advance the day on which Salcobrand would raise the price of each medicine.

> *"contaban con información previa sobre el día en que Salcobrand iba a alzar el precio de cada medicamento".*
>
> (TDLC, 2012, consid. 190)

(m) Punishment: the threat was a return to the price war. In the FNE's account, the two monitoring channels made the punishment threat credible. The punishment was a return to the price war, which had harmed all three chains.

> *"El funcionamiento eficiente de estos mecanismos de control hizo posible el seguimiento del acuerdo y, por ende, creíble la amenaza de castigo, bastante fuerte, pues consistía, naturalmente, en regresar al estado de guerra, tan perjudicial para todas las requeridas."*
>
> (FNE, 2008, para. 96, pp. 35–36)

(n) The agreement expanded after successful early batches. The FNE describes execution beginning in December 2007 with 62 medicines. It reports that 70 more were added in January 2008, 31 in February, and another 40 in March as the agreement's success was verified.

> *"incrementándose en número de productos a medida que se verificaba el éxito del acuerdo".*
>
> (FNE, 2008, paras. 14 and 17, pp. 5–6; paras. 19 and 22, pp. 8–12)

The same 19 December email states the expansion rule at a point when the count was still small. It reported increases for five products from four laboratories and proposed repeating the procedure with more products and laboratories over the following weeks. The tribunal read the message as evidence that the mechanism was working and that the chains intended to extend it.

> *"Hasta el momento se ha logrado que con 4 laboratorios hayamos subido el precio de 5 de sus principales productos. Dados estos buenos resultados esperamos repetir el 'procedimiento' con más productos y más laboratorios, en el transcurso de las próximas semanas."*
>
> (TDLC, 2012, consid. 95, 97)

(o) Scope, chronology, and outcome. The FNE's requests for information reached the chains at the end of March 2008 (TDLC, 2012, consid. 88). Its Oficio No. 419 to Salcobrand, in investigation Rol No. 1129-08, is dated 3 April 2008 (TDLC, 2012, consid. 102). The FNE filed its complaint on 9 December 2008. The tribunal approved FASA and the FNE's settlement on 13 April 2009 (TDLC, 2012, *Vistos* 4, p. 23): FASA admitted the participation quoted in (i), paid 1,350 annual tax units (*unidad tributaria anual*, UTA) to the Treasury, supplied a chronology of the price movements of the medicines in the complaint, and undertook to hand over further evidence (TDLC, 2012, consid. 12). The tribunal found coordination involving at least 206 medicines over a period covering at least December 2007 through March 2008 (TDLC, 2012, consid. 191). On 31 January 2012, it gave judgment against Cruz Verde and Salcobrand and fined each 20,000 UTA, the statutory maximum, which it noted came to about 3.4% of the two firms' combined sales of medicines for human consumption (the Farma category) in the year before the infringement (TDLC, 2012, consid. 205–206). The Supreme Court affirmed on 7 September 2012 (Corte Suprema de Chile, 2012).

(p) Communication after the investigation began. Asked about the effect of the investigation on the market, an executive of Laboratorio Garden House told the FNE that

Ahumada and Cruz Verde had stopped answering emails and calls and communicated only in person, while Salcobrand continued to use its normal communication channels.

> *"principalmente Ahumada y Cruz Verde han dejado de responder los mails y contestar llamadas, y ahora sólo se comunican personalmente. Salcobrand ha seguido con sus canales de comunicación normales".*
>
> (TDLC, 2012, consid. 88)

(q) The alternative explanations on the record. The respondents acknowledged the parallel movement of prices and offered noncollusive explanations for its cause. The tribunal recorded that the parallel conduct was fully established and acknowledged by both respondents.

> *"el que es además reconocido por las dos requeridas".*
>
> (TDLC, 2012, consid. 166)

Salcobrand's commercial manager, explaining his own email, said that what looked like an offer was only information passed to the laboratories. He said that he raised prices hoping that the others would do the same and applied the pricing policy after three weeks. He described the procedure as departing from the standing pricing rule, informing the laboratory, and conducting a competitor price check three weeks later.

> *"Yo subía con la esperanza que los demás lo subieran y a las tres semanas aplicaba política."*
>
> (TDLC, 2012, consid. 98)

Cruz Verde's defence had the same form. It claimed to have raised prices unilaterally in response to Salcobrand's increases, expecting FASA to do the same because FASA also had an incentive to raise.

> *"esperando que FASA hiciera lo mismo, ya que también tenía incentivos para ello".*
>
> (TDLC, 2012, consid. 165)

One recipient of the Salcobrand email understood the proposed change in pricing policy in its fourth point. The offer of specific days and the word "achieved" remained unexplained. The tribunal rejected these explanations.

> *"respecto del ofrecimiento de días específicos, no sabía por qué era y que tampoco sabía por qué el texto decía 'logrado'"; "la otra parte no la entiendo".*
>
> (TDLC, 2012, consid. 101)

Several executives said that "coordinated" in their emails meant telling a laboratory the date of an increase. The tribunal applied the verb's ordinary meaning: coordinating arranges means or efforts toward a common action, whereas communication alone falls outside that definition.

> *"concertar medios, esfuerzos, etc., para una acción común"; "y no 'comunicar algo'."*
>
> (TDLC, 2012, consid. 99–100)

The tribunal accepted that the expert report showed significant statistical anomalies. It also held that those anomalies could arise if Salcobrand became a price leader and FASA and Cruz Verde became followers. The statistical evidence therefore had to be combined with additional evidence to distinguish price leadership from collusion.

> *"también podrían explicarse como resultado de un escenario en que SB pasa a tomar el rol de líder de precios, y en que FASA y CV pasan a ser seguidores de SB"; "deben ser complementadas con evidencia adicional, para poder distinguir claramente entre ambas hipótesis".*
>
> (TDLC, 2012, consid. 159)

After considering the documentary and testimonial evidence, the tribunal rejected the alternative explanations (TDLC, 2012, consid. 132, 161, 192).

B Data and evidence from the price-event panel

This appendix has two tasks. First, it explains how I construct the daily panel, price tiers, and price-event panel used in Sections 2 and 3. Second, it reports additional evidence on the cost of below-cost pricing and on how the chains organized price increases. Appendix B.1 describes the data and analysis samples. Appendix B.2 explains how I construct the price tiers. Appendix B.3 measures the direct loss on medicines sold below cost. Appendix B.4 reports the organization and timing of price increases.

B.1 Data construction

The processed daily panel covers 222 medicines at Cruz Verde, FASA, and Salcobrand from 2006 through 2008. Each chain–medicine–day cell contains units sold and a nationwide revenue-weighted average transaction price in Chilean pesos (CLP). Separate daily listed-price series cover the same chain–medicine panel. A day with no transaction has no observed transaction price.

I use daily prices as recorded to examine how long price cuts lasted. For demand estimation, I keep only days when all three chains report positive sales and valid prices for the same medicine. I do not fill missing prices or quantities or apply a centered moving average.

Product attributes include the laboratory, active ingredient, therapeutic class, package size, pill count, brand or generic status, and prescription status. I compiled these attributes from a product catalog in the case records, DrugBank, and information I scraped from the Chilean online pharmacy `https://pharmazon.cl`. When the sources conflicted, I used the information from pharmazon.cl. Wholesale costs are Salcobrand's wholesale prices for November 2007 through May 2008, submitted as expert evidence to the Competition Tribunal (TDLC, 2012). Section B.3 explains how these costs, in CLP per unit, enter the below-cost loss calculation, including for earlier dates.

The data cover 222 medicines, and I solve the dynamic game for all 222. Of these, 210 completed their first successful coordinated price increase out of Tier 0 (defined in Section B.2) by 30 June 2008, using the three-chain success criterion in Section 2.4, part (ii), of the main paper. The regressions in Section 3 that explain the order of successful price increases use these 210 medicines.

B.2 Construction of price tiers

I distinguish three price tiers: below cost (Tier 0), margin restoration (Tier 1), and rent extraction (Tier 2). For each medicine, I calculate one benchmark price per tier, common across the three chains. I first average the three chains' weekly transaction prices, in thousands of Chilean pesos, and then average these weekly prices over each tier's window. Tier 0 uses the average price during the price-war period. Tiers 1 and 2 use average prices over stable periods beginning two weeks after the corresponding first successful increase.[2] When a stable Tier 1 price is unavailable, I multiply the medicine's Tier 0 price by the median Tier 1-to-Tier 0 price ratio among medicines with observed prices in both tiers and a price increase greater than 0.5 percent.

For Tier 2, I retain the observed stable prices for 56 medicines and predict the missing prices for the remaining 166 medicines. Of the 70 medicines that reached Tier 2 by 5 April 2008, 14 first did so in the final two weeks. Their stable-price windows would begin after the cutoff, so these 14 also require predicted benchmark prices despite having observed successful increases. Let p_{1j} and p_{2j} denote medicine j's Tier 1 and Tier 2 benchmark prices,

[2] Tier 0 averages weekly prices over 4 November–15 December 2007. The Tier 1 window begins two weeks after the first successful crossing of the Tier $0 \to 1$ boundary and ends before the first Tier $1 \to 2$ crossing or at the end of the week of 30 March–5 April 2008, whichever is earlier. The Tier 2 window begins two weeks after the first Tier $1 \to 2$ crossing and ends on 5 April 2008.

in thousands of Chilean pesos. On the 56 medicines with observed Tier 2 prices, I regress $\log(p_{2j}/p_{1j} - 1)$ on a constant, $\log p_{1j}$, an indicator for chronic-treatment medicines, and an indicator for prescription medicines. The Tier 1 benchmark follows the construction above, including imputation for three of these 56 medicines. Writing these four regressors as X_j and their ordinary-least-squares coefficients as $\widehat{\gamma}$, I set each missing Tier 2 price to

$$\widehat{p}_{2j} = p_{1j}\big[1 + \exp(X_j\widehat{\gamma})\big].$$

This prediction uses the fitted geometric proportional markup. I retain a minimum Tier 2-to-Tier 1 ratio of 1.005. This bound does not bind for any of the 166 predictions.

The imputation transports the observed-price relationship to medicines with missing Tier 2 prices. It does not correct for selection on unobserved determinants of coordination. The observed-price sample contains only one medicine not used for chronic treatment, and eight imputed medicines have Tier 1 prices outside its range. To compare the two prediction methods on the same scale, I evaluate their predictions of $\log(p_{2j}/p_{1j})$, transforming the regression prediction to $\log[1+\exp(X_j\widehat{\gamma})]$. Leaving out the laboratory of the only nonchronic medicine makes the regression unidentified. Across the other 55 medicines, leaving out one laboratory at a time when fitting both methods gives similar root-mean-squared prediction errors for the regression and median-ratio benchmark.

Table 1 counts successful increases by tier transition. A direct Tier $0 \to 2$ increase is one event that crosses both adjacent tier boundaries. Figure 2 illustrates the three price tiers for two medicines. The binary cartel indicator uses the existing manually coded chain-level coordination indicators and equals one when at least one chain is marked as participating. It does not distinguish the subsequent move to Tier 2.

Table 1: Successful price events by tier transition

	Count
Medicines in the processed panel	222
Successful Tier $0 \to 1$ events	206
Successful Tier $1 \to 2$ events	85
Successful direct Tier $0 \to 2$ events	7
All successful events	298

Note: The three successful-event rows are mutually exclusive and cover 31 December 2006–3 January 2009. Each event is one medicine-level increase attempt.

Figure 2: Price tiers and the binary cartel indicator

(a) Drug 7: two tier transitions versus one binary-indicator onset

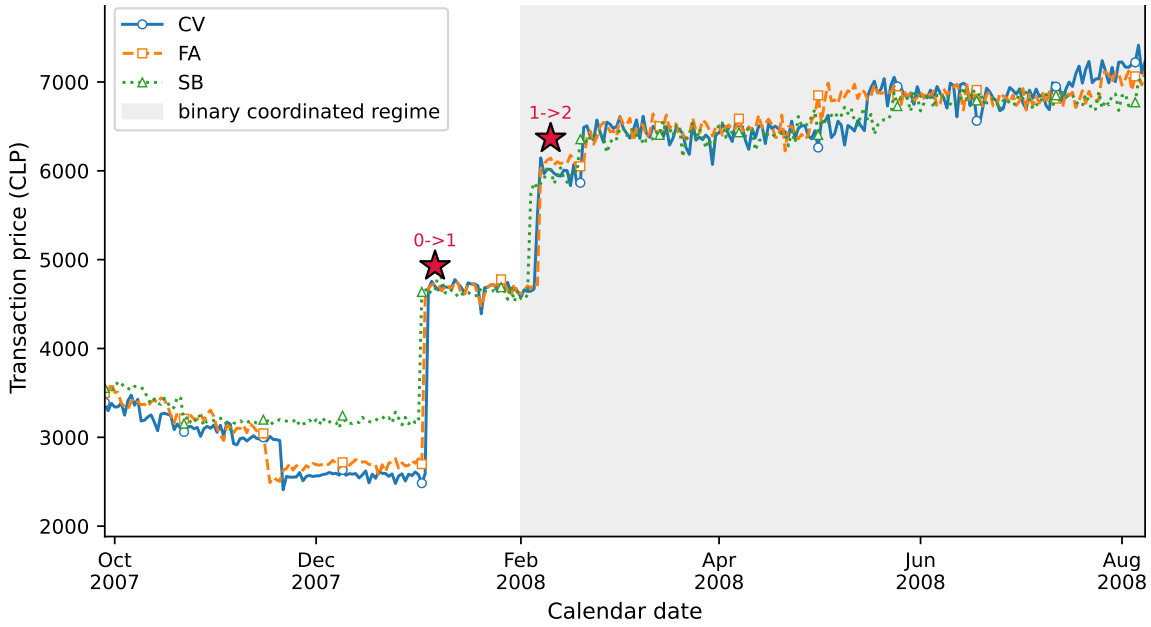

(b) Drug 10: two tier transitions versus one binary-indicator onset

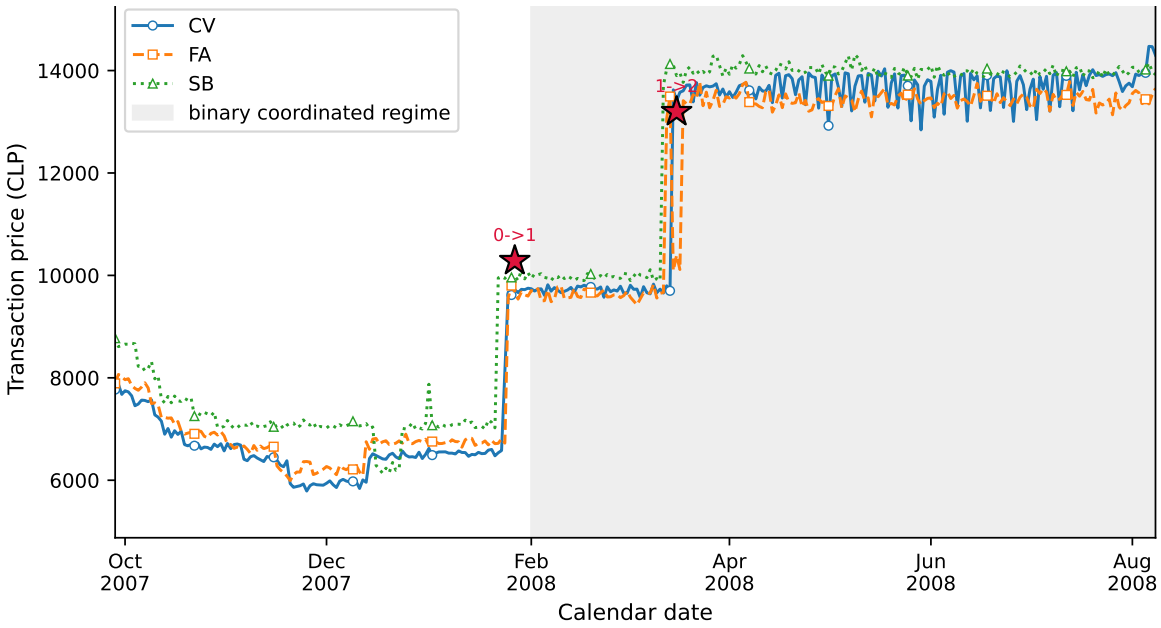

Note: The horizontal axes show calendar dates, and the vertical axes show transaction prices in CLP. Stars mark observed moves between price tiers. The gray area begins when at least one chain's indicator records entry into the coordinated-price regime. For display only, zero prices are treated as missing, isolated jumps larger than threefold are replaced by the median of adjacent values, and remaining gaps are filled from neighboring observations.

Observed price levels above Tier 2. Some medicines rise again after reaching Tier 2. I retain these later increases in Tier 2 when explaining discrepancies between observed paths and tier counts. They are not new moves between price tiers.

B.3 Monthly losses on below-cost medicine sales

This subsection measures the direct gross-margin loss from selling medicines below wholesale cost during the price war. For chain i and calendar month $\mathfrak{m}$, I calculate

$$\mathcal{L}_{i\mathfrak{m}}^{\text{money}} = \sum_{t \in \mathfrak{m}} \sum_{j} \max\{c_j - p_{ijt}, 0\} q_{ijt},$$

where p_{ijt} is the transaction price in CLP per unit and q_{ijt} is units sold by chain i for medicine j on day t, and c_j is medicine j's median observed Salcobrand wholesale cost from November 2007 through May 2008. The measure sums the peso shortfall on units sold below cost. It is not lost revenue. Because wholesale costs begin in November 2007, earlier losses use the later median cost and are approximate. One medicine with no observed cost receives the cross-medicine median. A missing daily transaction price contributes no measured shortfall.

Estimated monthly losses at Cruz Verde, FASA, and Salcobrand averaged 190, 124, and 78 million pesos in the first half of 2007 and reached respective late-2007 peaks of 910, 546, and 290 million pesos. Figure 3 shows that these losses fell as the coordinated increases restored positive margins.

Figure 3: Monthly losses on below-cost medicine sales

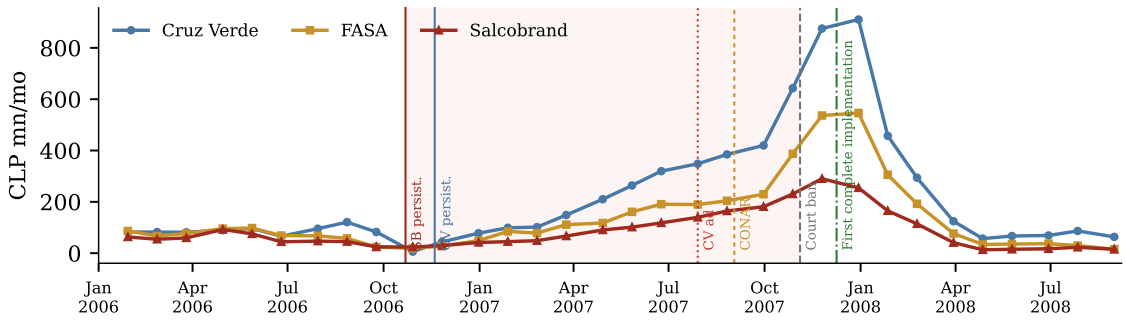

Note: Each series reports $\mathcal{L}_{i\mathfrak{m}}^{\text{money}}$ in millions of CLP. The vertical lines mark Salcobrand's and Cruz Verde's transitions to persistent cuts, the comparative-advertising campaign, the nonbinding advertising ruling, the binding advertisement ban, and the first successful coordinated increase (labeled "First complete implementation"). Values before November 2007 apply each medicine's November 2007–May 2008 median wholesale cost. "Court ban" in the figure denotes the advertisement ban.

B.4 How coordinated increases were organized

Laboratory batches. I group attempts to raise prices out of Tier 0 into a batch when the medicines come from the same laboratory and the attempts occur on the same date. A batch records observed attempts and may cover only part of a proposed medicine list (Appendix A, "(a) Sources and evidentiary conventions"). Before the first successful three-chain increase, none of the 42 attempts in 28 batches succeeded. Afterward, 206 of 216 attempts succeeded (Table 2). Batches became larger, and more laboratories participated. The dynamic model uses these patterns to specify when laboratories arrive and how many medicines they bring.

Table 2: Laboratory batches before and after the first successful increase

	Before first success	From first success onward	Full window
Dates	6 Nov.–8 Dec. 2007	9 Dec. 2007–23 Mar. 2008	6 Nov. 2007–23 Mar. 2008
Laboratory–date batches	28	83	111
Laboratories	13	36	36
First batches	13	23	36
Repeat batches	15	60	75
Medicine-level attempts	42	216	258
Successful Tier $0 \to 1$	0	203	203
Successful Tier $0 \to 2$	0	3	3
Reverted	23	3	26
Held without complete rival follow	19	7	26
Success rate (percent)	0.0	95.4	79.8
Mean attempts per batch	1.50	2.60	2.32
Unique medicines	39	209	210

Note: The two success rows are mutually exclusive. An unsuccessful attempt either reverted or remained at the higher price without complete rival follow. A first batch is the laboratory's first observed batch in this window. Later batches are repeats. Success rates refer to medicine-level attempts, not to the share of batches in which every medicine succeeded. A medicine can appear in both periods, so the unique-medicine columns are not additive.

Within-laboratory coordination order. Within each laboratory, chronic-treatment medicines completed their first successful increase out of Tier 0 earlier. Revenue, market size, and the other product attributes explain little of this order (Table 3). Both regressions have an adjusted R^2 below 0.07.

Table 3: Product attributes and order within laboratories

Variable	(1) Revenue	(2) Market size
Pre-ban log revenue	−0.041 (0.026)	—
Pre-ban log unit market size	—	−0.029 (0.035)
Chronic treatment	−0.216** (0.088)	−0.222** (0.087)
Prescription medicine	0.044 (0.094)	0.036 (0.094)
Generic medicine	0.004 (0.044)	0.004 (0.046)
log(1 + number of manufacturers)	0.036 (0.023)	0.031 (0.023)
Below-cost gap	0.006 (0.008)	0.001 (0.010)
Therapeutic-area fixed effects	Yes	Yes
Observations; laboratories	204; 30	204; 30
R^2	0.145	0.137
Adjusted R^2	0.067	0.058

Note: The outcome ranks each medicine's first successful event out of Tier 0 from earliest (0) to latest (1) within its laboratory. Ties receive their average rank. Revenue and unit market size are pre-ban daily means of the totals across the three chains. Log-transformed regressors use $\log(1 + x)$. The below-cost gap is the pre-ban mean of $\max\{c_j - p_{ijt}, 0\}$ across chain–medicine–days. Continuous variables are standardized before removing laboratory means; indicator variables are demeaned without rescaling. The regressions include broad therapeutic-area fixed effects, and standard errors are clustered by laboratory. Six singleton laboratories are excluded from the 210-medicine timing sample. ** denotes significance at the 5 percent level.

Leadership by chain and period. Salcobrand was usually the first chain to raise its price. It initiated 208 of 273 successful events during the cartel wave and 217 of 298 over the full sample (Table 4, Panel A).

Table 4, Panel B, compares how quickly medicines moved from Tier 0 to Tier 1 and from Tier 1 to Tier 2. For each type of increase, I use the 65 days starting with its first success. The number of medicines still waiting to make each move changes over time. To account for this, I divide the number of first successful moves by the total time medicines were able to make them.

On each day in the window, I count how many medicines could still make the increase. For Tier $0 \to 1$, these are medicines that have not yet left Tier 0. For Tier $1 \to 2$, these are medicines that reached Tier 1 before that day but have not yet reached Tier 2. A medicine is counted on the day it makes the increase. It starts contributing time in Tier 1 the next day. I add these daily counts and divide by seven to get medicine-weeks. The rate in Panel B is

$$\text{Moves per 100 medicine-weeks} = 100 \times \frac{\text{Number of first successful moves}}{\text{Total medicine-days}/7}.$$

It measures how often medicines moved per unit of time they could make the move. It is not the probability of success in a week.

Table 4: Leadership and tier-transition rates in matched windows

Panel A. Successful events initiated by each chain

	Cruz Verde	FASA	Salcobrand
Cartel wave	13	52	208
After 31 March 2008	11	5	9
Tier 0 → 1 events	12	39	155
Tier 1 → 2 events	11	15	59
Direct Tier 0 → 2 events	1	3	3
All successful events	**24**	**57**	**217**
Share of all 298 events (percent)	8.1	19.1	72.8

Panel B. Tier-transition rates during matched 65-day windows

Transition	Calendar window	First transitions	Medicine-days	Moves per 100 medicine-weeks
Tier 0 → 1	9 Dec. 2007–11 Feb. 2008	138	9,823	9.83
Tier 1 → 2	27 Jan.–31 Mar. 2008	66	8,577	5.39

Note: Panel A records who raised first, not who proposed the procedure. The cartel wave is 9 December 2007–31 March 2008 . Panel A's three transition rows cover 31 December 2006–3 January 2009 and do not overlap. In Panel B, later price cuts do not restart the count. A direct Tier 0 → 2 move ends the medicine's time in Tier 0 without entering the Tier 1 group and is excluded from both numerators. One such move occurs in the first window and two in the second.

After the investigation. Price-increase attempts continued after the investigation began, but the analysis-eligible event sample records 25 successful coordination events from 1 April through 31 December 2008, compared with 273 during the cartel wave (Table 4, Panel A). From 1 April through 31 December 2008, 115 attempts failed: 67 reverted, while 48 stayed at the higher price without both rivals following. Salcobrand initiated 82 of these attempts. Over the same period, there were also 49 sustained cuts from an established coordinated price.

There were 19 successful Tier 1 $\rightarrow$ 2 increases after 31 March 2008, excluding direct Tier 0 $\rightarrow$ 2 increases. Three involved Pfizer medicines on 25 July 2008, so medicines from the same laboratory were still sometimes raised together.

C Demand validation and payoff thresholds

This appendix reports within-week profit thresholds . It checks the sensitivity of low-price coefficients and model-action quantity effects to the demand specification . It compares observed quantity changes with fitted outcomes around actual price changes.

C.1 Daily quantity and profit

I calculate how a price change affects the chain's profit over the seven days it remains in place. For each medicine, one chain changes its price while the other two keep theirs unchanged. I compare this with all three chains keeping their original prices, and repeat the calculation for each chain. This first comparison covers the price-change week. The event-profit calculation in Figure 5 and Equation (2) also includes the following seven days, with that week's profit change discounted.

Medicine profit is the price minus unit cost, multiplied by quantity sold. A price change affects both the profit per unit and the number of units sold . Unit cost stays fixed. Let $\Delta\pi^{\text{medicine}}$ be the resulting change in daily medicine profit and Δq the change in daily quantity. With basket profit μ per unit of medicine sold, the profit change over the seven days is

$$\Delta\Pi^{\text{flow}}(\mu) = 7\{\Delta\pi^{\text{medicine}} + \mu\Delta q\}.$$

The daily changes are constant over these seven days in this calculation. Medicine profit is in thousand CLP and quantity in units, so μ is in thousand CLP per unit of medicine sold.

I set this profit change equal to zero to find the break-even basket profit:

$$\mu^* = -\frac{\Delta\pi^{\text{medicine}}}{\Delta q}.$$

The factor of seven cancels, so the cutoff is the same for daily and seven-day profit. For a cut that increases quantity, profit rises when $\mu > \mu_C^*$. For a solo increase that reduces quantity, it rises when $\mu < \mu_R^*$. These cutoffs concern profit during the price-change week . They do not by themselves determine the best action in the dynamic game. Table 5 compares the daily quantity changes and cutoffs before and after the advertisement ban.

Table 5: Predicted quantity changes and break-even basket profit

Action	Transition	Quantity change (%)		μ^* (thousand CLP)	
		Before	After	Before	After
Solo increase	Observed initial state to Tier 1	−23.8	−19.1	3.56	5.13
Solo increase	Tier 1 to Tier 2	−24.9	−21.8	3.40	4.53
Cut	Initial to price war	4.6	3.1	12.84	26.13
Cut	Tier 1 to price war	25.6	9.2	6.60	18.95
Cut	Tier 2 to Tier 1	26.9	10.0	4.51	16.40
Cut	First deeper price-war cut (5%)	4.5	3.0	13.32	26.25

Notes. Each row begins with 222 medicines and three chains (666 medicine–chain nodes); rivals' prices stay fixed. Initial prices are chain-specific observed prices. Cases without an actual price change in the stated direction are excluded from both reported medians. The valid pre/post node counts are 467/467, 666/666, 666/663, 527/527, 666/666, and 665/661 in row order. Basket profit is in thousand CLP per medicine unit. Cutoffs do not account for demand-estimation uncertainty.

For a solo increase from a chain's observed initial price to Tier 1, with both rivals keeping their initial prices, the median basket-profit cutoff rises from 3.56 before the advertisement ban to 5.13 afterward. The increase reduces profit when μ exceeds its cutoff, so a higher cutoff allows a larger basket profit before the increase becomes costly. For a cut from a chain's observed initial price to its price-war price, again holding rivals' prices fixed, the median cutoff rises from 12.84 to 26.13. The cut raises profit when μ exceeds its cutoff, so this type of cut is harder to make profitable after the ban.

The remaining rows apply the same within-week calculation at later price states. The Tier 1 → 2 row evaluates an unmatched increase out of the restored-margin state. For cuts, only the cutting chain changes price during the action week: it moves from Tier 1 to the initial price-war price, from Tier 2 to Tier 1, or from $0^{(0)}$ to $0^{(1)}$ in the first deeper five-percent price-war cut. In the dynamic game, a verified cut can move all three chains to the price-war punishment at the next week's opening; that continuation is deliberately excluded here. Within every cut row, the post-ban quantity gain is smaller and the break-even cutoff is higher. The intended comparison is before versus after the ban within a row, because starting prices and the size of the price change differ across rows.

C.2 Demand specifications and quantity effects of model actions

I vary the price-history controls to assess the sensitivity of the estimated low-price coefficient and the implied quantity effects, keeping the low-price-status definition fixed. Tables 6 and 7 use the same nine specifications, organized into four groups. The four groups change the price-history terms in demand Equation (1), $\gamma_{\text{inc}} A^{+}_{ijt} + \gamma_{\text{cut}} A^{-}_{ijt} + \gamma_H H_{ijt}$. Here A^{+} and A^{-} measure recent price increases and cuts, and H measures the gap above both rivals after a recent increase. I re-estimate demand for each row.

Group A changes the number of days summed in A^{+} and A^{-}. The headline model sums seven days. The alternative sums three. This also changes how long an increase is recent for H.

Group B replaces the sizes of price increases and cuts with indicators for whether any increase or cut occurred in the seven-day window, replaces H with an indicator for whether its conditions hold, or makes both replacements.

Group C retains the sizes but weights a change ℓ days ago by $2^{-\ell/T}$ in A^{+} and A^{-}. Here T equals one or three days. A change receives half as much weight every T days. The high-price term H is unchanged.

Group D changes only the price-gap condition for H. Instead of requiring both rivals to be at least five percent cheaper, it requires any strictly higher price or a ten-percent gap. The recent-increase requirement and the scale of the gap stay fixed.

Table 6: Low-price demand effects across specifications

Demand specification	Pre-ban	Post-ban
A. How many days of price history?		
Headline model	0.207 [0.179, 0.232]	0.062 [0.035, 0.084]
Three days	0.197 [0.166, 0.224]	0.058 [0.030, 0.082]
B. Amounts or yes/no indicators?		
Price changes: yes/no	0.221 [0.191, 0.247]	0.056 [0.028, 0.080]
High-price position: yes/no	0.203 [0.174, 0.229]	0.058 [0.031, 0.080]
Both: yes/no	0.217 [0.188, 0.243]	0.053 [0.025, 0.076]
C. Less weight on older price changes		
One-day half-life	0.178 [0.143, 0.210]	0.048 [0.020, 0.074]
Three-day half-life	0.183 [0.150, 0.214]	0.052 [0.024, 0.076]
D. When does the high-price effect apply?		
Any price above both rivals	0.201 [0.172, 0.226]	0.056 [0.029, 0.079]
Both rivals at least 10% cheaper	0.211 [0.183, 0.236]	0.067 [0.039, 0.089]

Notes. Entries are coefficients on low-price status in log share relative to the outside good. The low-price rule remains at least five percent below the highest chain price. All specifications are re-estimated on the headline sample of 589,624 IV observations and 222 medicines, with chain–medicine and date fixed effects. As in the headline estimation, an unavailable calendar-day price change is set to zero. Brackets give 95-percent intervals from the same 199 medicine bootstrap draws.

Across all nine specifications, Table 6 shows a smaller estimated low-price coefficient after the advertisement ban.

In the full headline sample of 222 medicines (666 medicine–chain nodes), the standardized five-percent cut has a two-week break-even basket profit of 12.73 thousand CLP before the ban and 25.86 afterward. Table 7 extends the quantity calculation in Table 5 to include the following week's price-history effects. All nine specifications use the headline sample of 222 medicines and the same 589,624 IV observations. Holding this sample fixed makes the specifications comparable. It uses the same five-percent initial cut as Table 5, but adds the following week's price-history correction.

Table 7 uses the initial increase and cut in Figure 5 and Section 5.1. One chain raises from its observed initial price to Tier 1, provided this is an increase, or makes the first ordinary five-percent cut from its own initial price. Rivals hold their prices for seven days. Next week the unmatched raiser returns to its initial price; after a cut all chains move to 95 percent of their own initial prices. I add the action-week quantity change and the next-week price-history correction, using the same two-week baseline as Figure 5.

The table asks whether the direction of the pre-ban/post-ban difference in quantity effects survives changes in the demand specification. All specifications, including the headline

model, are re-estimated on the same sample as the headline demand equation. Unavailable calendar-day price changes are set to zero, and the chain–medicine effects used for action payoffs are recovered from the full outcome panel. These calculations exclude continuation values, including the price-war punishment after verification.

Table 7: Event quantity effects of the initial increase and cut

	Increase: quantity loss (%)		Cut: quantity gain (%)	
Demand specification	Pre-ban	Post-ban	Pre-ban	Post-ban
A. How many days of price history?				
Headline model	12.17 [10.68, 14.05]	9.82 [8.65, 11.73]	2.30 [1.90, 3.39]	1.54 [1.10, 2.20]
Three days	7.79 [6.99, 9.06]	6.12 [5.36, 6.95]	2.42 [2.04, 3.50]	1.72 [1.30, 2.36]
B. Amounts or yes/no indicators?				
Price changes: yes/no	1.46 [−2.12, 4.93]	−0.96 [−4.89, 2.22]	8.81 [7.09, 11.39]	8.06 [6.04, 10.03]
High-price position: yes/no	11.53 [10.48, 12.72]	9.86 [8.85, 11.07]	2.33 [1.93, 3.43]	1.58 [1.13, 2.25]
Both: yes/no	0.32 [−3.66, 3.79]	−1.39 [−5.64, 1.99]	8.84 [7.07, 11.40]	8.05 [6.06, 9.99]
C. Less weight on older price changes				
One-day half-life	10.12 [8.89, 11.68]	8.19 [7.07, 9.75]	2.52 [2.10, 3.59]	1.87 [1.35, 2.46]
Three-day half-life	10.65 [9.22, 12.17]	8.63 [7.51, 10.21]	2.70 [2.26, 3.82]	2.01 [1.51, 2.67]
D. When does the high-price effect apply?				
Any price above both rivals	12.45 [10.98, 14.43]	9.96 [8.81, 11.92]	2.26 [1.86, 3.34]	1.50 [1.05, 2.16]
Both rivals at least 10% cheaper	11.95 [10.67, 13.80]	9.51 [8.33, 11.54]	2.33 [1.94, 3.45]	1.56 [1.13, 2.22]

Notes. Medians over 467 eligible initial-to-Tier 1 increases and 666 initial cuts (222 medicines). Event quantity changes add seven action days and seven next-week history-correction days; the denominator sums no-history quantities at each week's price vector. Counterfactual paths begin with no prior price-change exposure. All models are re-estimated on the same 589,624 headline IV observations; unavailable calendar-day price changes are set to zero. Brackets give 95-percent intervals from 199 common medicine bootstrap draws.

Table 8 reports the corresponding break-even basket profit for the standardized five-percent cut. Every specification implies a higher cutoff after the advertisement ban. The cutoff levels are substantially smaller when recent price changes enter only as indicators. Among specifications that retain change magnitudes, the pre-ban cutoff ranges from 8.14 to 12.98 thousand CLP. The post-ban cutoff ranges from 13.27 to 27.14 thousand CLP.

Table 8: Break-even basket profit for a five-percent cut across demand specifications

	μ_C^* (thousand CLP)	
Demand specification	Pre-ban	Post-ban
A. How many days of price history?		
Headline model	12.73 [10.46, 14.74]	25.85 [19.55, 34.76]
Three days	11.60 [8.87, 13.59]	19.38 [15.82, 24.89]
B. Amounts or yes/no indicators?		
Price changes: yes/no	1.05 [0.75, 1.65]	1.65 [1.22, 2.37]
High-price position: yes/no	12.47 [10.25, 14.57]	25.59 [18.90, 33.57]
Both: yes/no	1.05 [0.76, 1.64]	1.67 [1.24, 2.36]
C. Less weight on older price changes		
One-day half-life	9.74 [7.86, 11.49]	16.06 [12.59, 20.17]
Three-day half-life	8.14 [6.29, 10.23]	13.27 [9.97, 17.29]
D. When does the high-price effect apply?		
Any price above both rivals	12.98 [10.79, 15.21]	27.14 [20.52, 36.89]
Both rivals at least 10% cheaper	12.49 [10.21, 14.45]	24.87 [18.57, 32.70]

Notes. Each cutoff sets the median profit change across 666 medicine–chain nodes to zero. Profit includes the seven-day action week and the discounted seven-day next-week price-history correction . Continuation values are excluded. A cut is profitable for $\mu > \mu_C^*$. Brackets give 95-percent intervals from the same 199 medicine bootstrap draws.

In every specification, the point estimates imply a smaller initial-increase quantity loss and a smaller initial-cut gain under post-ban demand. Their magnitudes and precision depend on the price-history specification. These comparisons do not establish that the ban caused the differences.

C.3 Observed outcomes around price changes

Realized losses after unmatched increases. For unmatched increases, I use the seven attempts with complete data for seven days starting on the increase date. They involve six medicines, all with Salcobrand raising first. For each day, I compare its actual quantity with a benchmark based on its rivals' sales, adjusted for its usual quantity relative to theirs. I add both quantities over the week and calculate the percentage shortfall. I then compare this observed loss with the predicted loss for Salcobrand and the same medicines under a Tier 1 $\rightarrow$ 2 increase. Salcobrand raises to Tier 2 for seven days while both rivals stay at Tier 1. These observed attempts involve Tier 1 $\rightarrow$ 2, whereas Appendix C.2 considers

increases from the observed initial prices to Tier 1.

Table 9, Panel A, groups the predictions by the same four specification changes as Table 7. It reports an observed median loss of 20.13 percent and a predicted median of 21.99 percent under the main specification. Replacing the sizes of recent increases and cuts with yes/no indicators predicts much smaller losses. Retaining the sizes of price changes therefore comes closer to the observed median.

This agreement concerns the median, not every event. The main specification's median absolute prediction error is 6.47 percentage points. Actual prices can reverse, and rivals can change their prices. The prediction holds the price path fixed. The sample is small. One medicine has overlapping attempts, and these observations also enter demand estimation. This is a check on the approximate size of the loss. It is not independent validation or a basis for choosing the specification.

Realized gains after cuts. Panel B reports actual seven-day quantity gains and fitted gains for the same cut events. Both use the same rival-adjusted quantity benchmark as Panel A.

Table 9: Realized and predicted seven-day quantity changes

Panel A. Losses after unmatched increases

Observed benchmark or specification	Median loss (%)
Observed: seven failed attempts	20.13
1. How many days of price history?	
Headline model	21.99
Three days	13.14
2. Amounts or yes/no indicators?	
Price changes: yes/no	15.43
High-price position: yes/no	19.81
Both: yes/no	12.43
3. Less weight on older price changes	
One-day half-life	21.12
Three-day half-life	23.77
4. When does the high-price effect apply?	
Any price above both rivals	22.51
Both rivals at least 10% cheaper	21.74

Panel B. Gains after actual cuts (%)

	Pre-ban	Post-ban
Observed median gain	0.39	12.20
Predicted median gain	0.45	8.77

Notes. Panel A weights seven attempts equally. The observed quantity benchmark uses rivals' current sales and the mean own-to-rival ratio on four preceding matching weekdays. Predictions use the fixed Tier $1 \to 2$ path. No uncertainty intervals are reported for this small sample. Panel B uses 2,845 pre-ban and 163 post-ban cuts with complete inputs. Gain is $100(Q/Q^{\text{benchmark}} - 1)$ using seven-day totals, with the same benchmark construction applied to actual and fitted quantities. Predictions use recorded prices. This is an in-sample comparison.

Panel B predicts positive median quantity gains in both periods, close to the observed median before the ban and below it afterward. The columns compare different events, cut sizes and rival responses, rather than the same model action across demand regimes. Cutting incentives also depend on medicine margins and the price-war continuation after verification (Section 5.2).

D Computational implementation

This appendix explains how I turn the model in Section 5 into simulated price paths and parameter estimates. The four steps are to prepare the inputs, solve firms' choices, simulate the weekly path, and compare it with the data.

D.1 Inputs, timing, and state variables

I use the estimated demand coefficients, the drug–chain price and demand primitives, and the event records.

The three blocks are the price war in weeks 52–82, the advertising campaign in weeks 83–95, and the laboratory period in weeks 96–117. The first two use pre-ban demand and the last uses post-ban demand. Prices carry across the boundaries. The advertisement ban took effect on 6 November 2007, within model week 96 (4–10 November). The model applies post-ban demand throughout that week. The demand estimation sample instead starts on 12 November, after the transition period.

A price decision occurs on the first day of a model week and applies for seven days. Its opening-day price change therefore stays in the seven-day demand history throughout the action week. The payoff correction in Section 5.1 accounts for next week's price history when prices are carried forward, restored, or synchronized. The price-history variables A^+, A^-, and H are not separate state variables.

The price states and public information are those defined in Section 5.2. The public state records verification (R_t), completed rent-extraction attempts (S_t^r), and incomplete attempts (F_t^r). These counts determine the leader's subjective weight m_t^r through Equation (8). All medicines use the same opening information. Outcomes during the week change next week's information.

D.2 Solution algorithm

I solve each block as an infinite-horizon stationary game. Firms expect its demand and review process to continue, rather than anticipating the next block. I also solve the decentralized game that unverified followers believe they are playing in Game 2. It has the campaign block's structure and cut depths, but post-ban demand. Salcobrand takes those followers' choices into account until verification changes their information.

Given continuation values, I first solve the increase stage. Choices are simultaneous in Game 1. In Game 2, followers use the decentralized policies before verification and choose sequentially afterward. I use the resulting value as the value of all three chains holding at

the preceding cut stage, and solve their cut-or-hold choices. I then average over no review, a cut, and an increase after all chains hold, using Equation (13). These updated values become the continuation values in the next iteration. The review probability enters this calculation once. The action scale τ is not estimated. I set it to the median positive absolute one-period payoff gap across the cut and increase comparisons at $\mu = 0$, divided by $2 \log 9$. The reference gap therefore implies a logit choice probability of 81/82. The same fixed scale applies to cut-or-hold, lead-or-wait, and follower decisions in both games.

Learning requires one further approximation. The same subjective weight can come from a short or a long history of attempts, but one new outcome changes it more after a short history. Instead of keeping both counts as separate dimensions of the value calculation, I use a grid of weights m. At each weight, I average over the simulated distribution of the effective count $\nu^r + S_t^r + F_t^r$. This gives probabilities of moving between grid values next week. Section D.4 explains how this transition is prepared and held fixed during estimation.

D.3 Forward simulation

Each simulated path follows the calendar through week 117. I keep each medicine's price state, three chain prices, punishment status, and eligibility for review, together with the common public information. Temporary price differences are kept in the simulation even though they are not additional points on the value-function grid. The simulation updates the actual counts of completed and incomplete attempts and uses the policy at the nearest weight on that grid.

Review opportunities. Game 1 reviews each eligible medicine with probability ω_1 before the campaign and ω_2 during it. Game 2 uses observed laboratory dates and medicine counts. At each date, I randomly sample eligible medicines from that laboratory without replacement, up to the observed count. Unpunished medicines below Tier 1 are eligible throughout Game 2. Unpunished Tier 1 medicines become eligible for rent-extraction reviews after verification.

Each review considers cuts before increases. Any cut cancels that week's increase opportunity. Before verification, the medicine remains eligible next week. Verification takes effect the following week. Reviews blocked by cuts during the verification week are retried then.

To fill gaps in randomly assigned laboratory reviews, I draw an independent deadline for each medicine and path when informed play begins, uniformly from that week through the penultimate week (the final week if it alone remains). From its deadline onward, an eligible, unpunished medicine below Tier 1 receives a supplementary review until it has a

review after verification without a cut, regardless of whether an increase succeeds. Cut-blocked supplements remain pending subject to continued eligibility. Supplements exclude Tier $1 \to 2$. Combined sources allow at most one review per medicine per week. These calendar rules govern simulation. Stationary values use the average review opportunity.

Weekly sequence. I start from observed week-52 prices and states, with no verification, learning outcomes, punishment, or temporary simulated increases. Each week then follows four steps.

1. **Open the week.** If the preceding week's increase was incomplete, restore participating chains' pre-attempt prices. Carry forward the price states and public information, and form the review set described above.

2. **Draw cuts.** Punished medicines follow their cut-only policy every week without a review. For each reviewed, unpunished medicine, draw the three chains' simultaneous cut-or-hold choices from the solved policy. Before verification, I also impose the observed weekly total of post-ban follower cuts, sampling the affected drug–chain identities from the eligible simulated medicines. These totals are inputs, not fitted targets.

   Any cut ends that medicine's decisions for the week. Ordinary cuts move one step down the price-war ladder. Campaign and Game 2 cuts use the calibrated drug-specific steps. These same steps enter the unverified followers' game and later punishment cuts. A first cut after verification instead starts punishment at $0^{(0)}$ next week and removes future increase opportunities. After an unverified cut, the chains' prices synchronize at the lower price next week and the medicine remains eligible for review. Already-punished medicines remain in punishment.

3. **Draw increases if all chains held.** In Game 1, draw the three simultaneous increase-or-hold choices. In Game 2, draw the candidate leader's lead-or-wait choice and then the followers' choices. Salcobrand is the only candidate before verification. Afterward, the chains consider leading in a random priority order. The first willing chain leads. If all wait, no chain leads. Unverified followers use their decentralized-game policies. Verified followers choose sequentially in both specifications. Unit Confidence fixes the leader's subjective weight at one without forcing followers to join. Only three-chain participation advances the persistent price tier. An incomplete increase lasts for the action week and does not prevent a later review.

4. **Close the week.** After all medicines have acted, update prices, punishment, and public information. Each Game 2 week that opens unverified has one common verification draw

with probability η_0, regardless of whether an increase occurred. Verification changes play from the next week and initializes the Adaptive weight at m_0^r. Unit Confidence sets it to one. Verification does not complete current increases. Only initiated Tier 1 $\rightarrow$ 2 proposals after verification add completed or incomplete attempts to the learning counts. Punished medicines add none. Record prices, states, and the dates of successful tier crossings for the moments below.

D.4 Estimation and inference

Criterion. For each candidate parameter vector, I compare the data with averages from 100 simulated paths. I divide each discrepancy by its scale below, square it, and average over the relevant weeks or states. This produces nine losses, grouped as follows. The criterion gives each loss equal weight, as in Section 5.4. Screened observed cut events and all simulated cut actions have different definitions, so cut counts are excluded.

Cumulative tier-transition counts (two losses). I compare the counts for Tier $0 \rightarrow 1$ and Tier $1 \rightarrow 2$ in weeks 101, 104, 110, 116, and 117. The observed counts are $(2, 62, 139, 206, 208)$ and $(0, 1, 18, 61, 70)$, respectively. Each medicine counts once per boundary. The first count includes only successful laboratory-mediated increases in Game 2. Independent Game 1 increases do not enter it. Each loss uses a scale of 5 medicines and averages over the five weeks.

Weekly price paths (three losses). I compare weekly average prices separately in weeks 52–82, 83–95, and 96–117. In both data and simulations, I average daily prices within each drug–chain pair and week, then across pairs. Temporary price changes enter for their actual duration. Each loss uses a scale of 0.25 thousand CLP and averages over the weeks in its block.

Price states and the block boundary (three losses). The first loss compares mean price-war depth among medicines still in price-war states in week 117. Depth counts ordinary five-percent price steps from the initial price, so the first price-war state has depth one. The observed mean is 1.85 steps. The scale is 0.5 step.

The second compares the week-96 average price, using the same construction and scale as the price paths. Its observed value is 10.350 thousand CLP. Week 96 thus enters both this loss and the Game 2 price-path loss.

The third compares the end-of-week 96 simulated price-state shares, after the first Game 2 update, with observed week-96 prices mapped to the same grid. For observed share p_s in

state s, the scale is $\sigma_s = \sqrt{\max\{p_s(1-p_s), 1/222\}/222}$. I average the squared standardized differences across all grid states. 222 is the number of medicines.

Verification timing (one loss). For each week from 96 to 117, I compare the share of paths verified by that week's close with the observed timing benchmark. It is zero through week 100 and one from week 101 onward. This benchmark is the first observed complete laboratory-mediated increase, not a direct observation of firms' knowledge. I use a scale of 0.25 and average across the 22 weeks. Verification by a week's close means informed play begins the following week. Paths never verified remain in the denominator. Only the separately reported mean verification week excludes them.

Parameter search. Both models estimate $(\mu, \omega_1, \omega_2, \eta_0)$, with common basket profit across blocks. Adaptive Confidence also estimates m_0^r and ν^r. Basket profit is bounded by $\mu \in [0, 50]$ thousand CLP per customer-equivalent, meaning one unit of medicine demand. The three probabilities and initial weight each lie in $[10^{-4}, 1 - 10^{-4}]$. Prior strength lies in $[10, 1000]$.

I use coordinate pattern search. I try changes in each parameter and shrink the search steps when none improves the criterion. All candidates and the reported fit use 100 paths with the same random draws (seed 91003). Each candidate solves the games and simulates the full path.

For each candidate parameter vector, I compute approximate game solutions with a choice-probability tolerance of 10^{-6}. I reduce the update step for medicines whose choice probabilities repeatedly oscillate. The reported estimates use a policy-iteration cap of 30 and retain the current policy for any game that reaches the cap.

Convergence diagnostics. I assess each solved game using the number of policy iterations, the maximum change in choice probabilities, and whether it reaches the 10^{-6} tolerance. Estimation uses capped policy iteration and conditions on a simulation-based learning transition fixed before the parameter search.

For Adaptive Confidence, I update the learning transition in Section D.2 once using simulated paths at the starting parameters, assigning weight 0.25 to the simulated transition frequencies and 0.75 to the existing transition. I then estimate the parameters conditional on this updated transition, which remains fixed during the parameter search.

Inference. The completed 100-replication laboratory-block bootstrap samples laboratories with replacement, keep their medicines and price histories, and re-estimate demand and both

dynamic models on each draw. The main-text table reports its standard errors. Within each draw, I reconstruct the Adaptive learning transition with the same one-step update used for the headline estimate and then hold that transition fixed while estimating the six Adaptive parameters. The retained earlier batch held learning parameters fixed and is not used for the current Adaptive standard errors. The main-text table reports the current headline estimates. Detailed sensitivity results follow below.

E Additional dynamic-model results and sensitivity checks

E.1 Additional model fit

Table 10 supplements the price and cumulative-count paths in the main paper with the remaining baseline moments and all nine loss components. Both models overpredict terminal price-war depth and the Game 2 opening price. Adaptive Confidence has smaller losses for Tier 1 $\rightarrow$ 2 counts, terminal depth and verification timing, but a larger Game 2 price-path loss. The week-96 state-distribution loss remains substantial in both specifications. These comparisons describe fit at separately estimated parameters.

Figure 4 shows the full distributions behind two of these moments. Both models place too many medicines at shallow price-war depth one and too few at several intermediate depths at the Game 2 boundary. The verification curves are compared with the first-success timing benchmark used in estimation. The benchmark does not measure firms' information directly. Adaptive Confidence reaches full verification by the end of the simulation, while some Unit Confidence paths remain unverified.

Table 10: Additional baseline moments and SMM loss components

Panel A. Additional moments

Moment	Reference	Unit	Adaptive
Mean price-war depth, week 117 (steps)	1.850	2.802	2.288
Mean transaction price, week 96 (thousand CLP)	10.350	10.522	10.568
Verified by week 101 (%)	100	44.0	61.0
Verified by week 117 (%)	100	89.0	100.0

Panel B. Standardized loss components

Loss block	Unit	Adaptive
Cumulative Tier 0 $\rightarrow$ 1 counts	18.590	10.933
Cumulative Tier 1 $\rightarrow$ 2 counts	27.612	0.366
Prices before the campaign	1.141	1.161
Prices during the campaign	0.314	0.344
Prices in Game 2	2.648	7.272
Terminal mean price-war depth	3.627	0.766
Week-96 mean price	0.472	0.757
Week-96 state distribution	5.515	5.596
Verification timing	1.746	0.987
Overall objective	6.852	3.131

Notes. Unit and Adaptive use the current baseline estimates and 100 simulation paths. In Panel A, the first two reference values are constructed from observed prices using the model's state mapping and price aggregation. Price-war depth is the mean among medicines remaining in price-war states. One step is an ordinary five-percent cut. Week 96 is 4–10 November 2007 and week 117 is 30 March–5 April 2008. The final two reference values are the maintained verification-timing benchmark, not observed shares of informed firms. The model shares include all paths. Panel B reports mean squared standardized discrepancies, with scales defined in Appendix D.4. The objective is the equally weighted mean of the nine losses. These are in-sample fit diagnostics from approximate game solutions and a fixed, approximate learning transition.

Figure 4: Additional baseline state and verification fit

(a) Price-state distribution at the Game 2 boundary

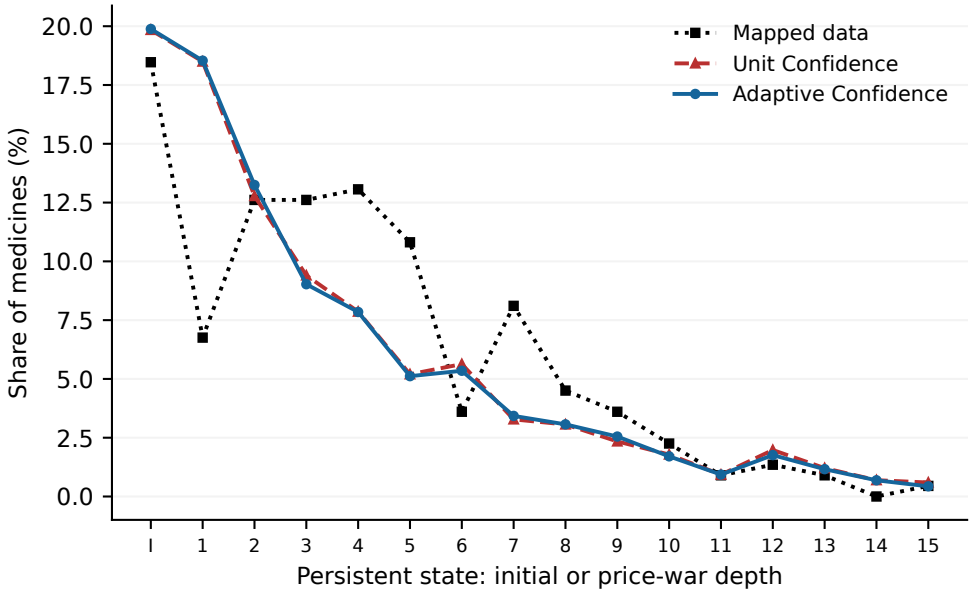

(b) Verification timing

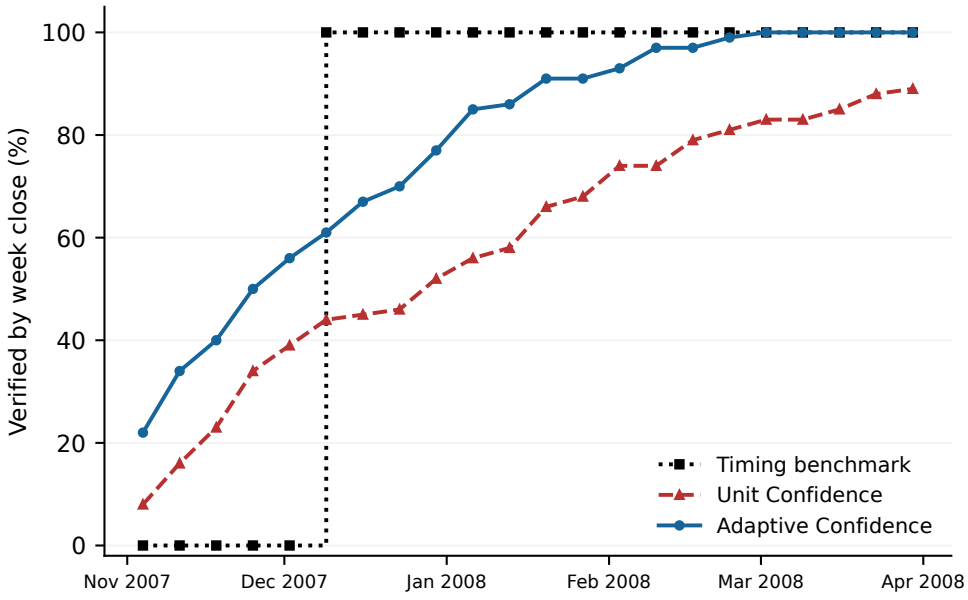

Notes. Panel (a) compares end-of-week 96 model state shares with observed week-96 prices mapped to the same grid. I denotes the initial state. Integer labels count ordinary five-percent price-war steps through the maximum observed week-96 depth. Coordinated tiers and deeper model-grid states are omitted from this boundary comparison. Panel (b) reports the share of all 100 paths verified by each week's close. The black timing benchmark is zero before the week beginning 9 December 2007 and one thereafter. It is not an observed information-state distribution. Axis dates identify the beginning of each week. Unit and Adaptive use their separately estimated baseline parameters and approximate game solutions.

E.2 Review-process sensitivity

Tables 11 and 12 compare alternative laboratory review processes and discount factors, respectively.

I compare the observed schedule with stochastic laboratory arrivals and adjacent-week swaps. Under stochastic arrivals, each laboratory's first opportunity follows its calibrated first-arrival probability. After the first opportunity, later weeks use its calibrated repeat-arrival probability. Conditional on arrival, the day is drawn uniformly within the week and the number of offered medicine-review slots is a positive-binomial draw bounded by the laboratory's medicine portfolio. The calibration matches the cumulative number of laboratories receiving a first opportunity and total offered slots in expectation. The adjacent-week design partitions the laboratory period into disjoint pairs of weeks and, independently for each laboratory and pair, exchanges the two complete observed batch schedules with probability one-half. It therefore preserves each laboratory's batches and total offered slots exactly while perturbing their timing. One offered slot is capacity to review one medicine. Used reviews depend on the eligible pool, and completed increases remain endogenous. Both alternatives retain the verified Tier $0 \to 1$ supplementary coverage reviews and verification-week retry used in the headline, with no supplementary Tier $1 \to 2$ review. Each row re-estimates four parameters for Unit Confidence and six for Adaptive Confidence.

Table 11: Sensitivity to the laboratory review process

Schedule	Specification	Objective	Tier $0 \to 1$	Tier $1 \to 2$
Observed schedule	Unit	11.644	199.73	129.37
Observed schedule	Adaptive	3.351	195.56	64.98
Adjacent-week swap	Unit	6.243	190.94	107.13
Adjacent-week swap	Adaptive	3.136	194.08	67.26
Stochastic arrival	Unit	5.233	192.78	98.82
Stochastic arrival	Adaptive	3.565	197.01	68.81

Notes. The tier columns are mean cumulative counts of distinct medicines crossing each boundary by the end of week 117, averaged over 100 simulated paths. Objective is the dimensionless mean of the nine standardized loss blocks defined in Section D.4. The observed-schedule rows give the benchmark for these comparisons. Compare schedules within each model.

Across all three review processes, Adaptive Confidence produces a lower objective and keeps the terminal Tier $1 \to 2$ count between 64.98 and 68.81, close to the observed count of 70. Unit Confidence remains more sensitive to the review process and produces between 98.82 and 129.37 Tier $1 \to 2$ transitions.

E.3 Discount-factor sensitivity

I compare annual discount factors 0.70 and 0.90 with the headline value 0.80, using $\beta = \beta_{\text{annual}}^{1/52}$ throughout each run. I re-estimate both specifications under the observed laboratory schedule, including the learning parameters for Adaptive Confidence, and keep the same coverage reviews, retry rule, calibrated cut depths and objective. I report the objective and terminal cumulative counts of distinct medicines crossing the Tier $0 \to 1$ and Tier $1 \to 2$ boundaries.

Table 12: Sensitivity to the annual discount factor

Annual factor	Specification	Objective	Tier $0 \to 1$	Tier $1 \to 2$
0.70	Unit	8.309	191.01	119.15
0.70	Adaptive	3.039	194.98	68.33
0.80	Unit	11.644	199.73	129.37
0.80	Adaptive	3.351	195.56	64.98
0.90	Unit	9.036	195.49	121.28
0.90	Adaptive	2.397	196.64	66.22

Notes. The tier columns are mean cumulative counts of distinct medicines crossing each boundary by the end of week 117, averaged over 100 simulated paths. Objective is the dimensionless mean of the nine standardized loss blocks defined in Section D.4. The 0.80 rows give the benchmark for these comparisons. Compare discount factors within each model.

Adaptive Confidence has the lower objective at all three discount factors, and its terminal Tier $1 \to 2$ count remains between 64.98 and 68.33. The contrast with Unit Confidence therefore does not depend on the headline annual discount factor of 0.80.

References

Corte Suprema de Chile. Sentencia rol no. 2578-2012: Farmacias Cruz Verde y Salcobrand con Fiscalía Nacional Económica. Sentencia Rol 2578-2012, Corte Suprema de Chile, Tercera Sala, 2012. URL https://www.fne.gob.cl/wp-content/uploads/2012/09/csse_03_2012.pdf. Judgment of 7 September 2012 affirming TDLC *Sentencia* No. 119/2012.

Farmacias Cruz Verde S.A. Contesta requerimiento, 2009. URL https://consultas.tdlc.cl/download/17d3464c3643f6226913c64bc2dae4e8. Party answer, TDLC Rol C-184-08, main docket, fojas 452–580.

Farmacias Cruz Verde S.A. Recurso de reclamación contra la sentencia no. 119/2012, 2012. URL https://consultas.tdlc.cl/download/7e89114d68d44055e24c75e3136fe34f. Party appeal, TDLC Rol C-184-08; testimony of Sergio Purcell reproduced on p. 110.

Fiscalía Nacional Económica. Observaciones a la prueba, 2011. URL https://consultas.tdlc.cl/download/8060d3720a3bf01bdca9cf152403afb9. 17 August 2011, TDLC Rol C-184-08, main docket, fojas 7037–7269.

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

Javier Núñez, Tomás Rau, and Jorge Rivera. Informe pericial sobre el requerimiento de la FNE en contra de farmacias ahumada s.a. y otros. Informe pericial, rol c no. 184-2008, Tribunal de Defensa de la Libre Competencia, Chile, 2010. URL https://consultas.tdlc.cl/download/94eda60da208cdc893a52eadb725c810. Court-appointed expert report dated 26 October 2010.

Servicio Nacional del Consumidor. Sernac denuncia a Cruz Verde por su publicidad "Desafío Cruz Verde". SERNAC press communication, 31 October 2007. URL: https://www.sernac.cl/portal/604/w3-article-918.html., 2007.

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.