Online Appendix to "Equilibrium Transition from Loss-Leader Competition: How Advertising Restrictions Facilitate Price Coordination in Chilean Pharmaceutical Retail"

Yu (Jasmine) Hao

July 1, 2026

This Online Appendix accompanies the main text and contains supplementary demand, event-coding, robustness, and mechanism checks. Cross-references of the form § n, Table n, and equation (n) refer to the main text; appendix material is numbered A–I below. 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 Demand-fit decomposition

Goodness-of-fit formulas. The within-nest share prediction used for the diagnostics of §4 is

$$\hat{s}_{i|jt} \;=\; \frac{\exp\big(\hat{u}_{ij(t)}/(1-\hat{\sigma})\big)}{\sum_{i'\in\mathcal{I}} \exp\big(\hat{u}_{i'j(t)}/(1-\hat{\sigma})\big)}, \qquad \hat{u}_{ij(t)} = \hat{\phi}_{ij} - \hat{\alpha}_{\text{eff}}(t)\, p_{ijt}, \tag{1}$$

with $\hat{\phi}_{ij}$ the within-(chain, drug) fixed effect and $\hat{\alpha}_{\text{eff}}(t) = \hat{\alpha}_0 + \hat{\alpha}_1 D_t$. Only the structural parameters and prices enter the right-hand side, so (1) is a non-circular prediction, not the mechanical contraction-mapping R^2 of Conlon and Gortmaker [2020, p. 32]. Own- and cross-price elasticities use the standard nested-logit formulas

$$\varepsilon^{\text{own}}_{ijj,t} = -\alpha\, p_{ijt}\left[\tfrac{1}{1-\sigma} - \tfrac{\sigma}{1-\sigma}\, s_{i|jt} - s_{ijt}\right], \tag{2}$$

$$\varepsilon^{\text{cross}}_{ijk,t} = \alpha\, p_{kjt}\, s_{kjt}\left[1 + \tfrac{\sigma}{(1-\sigma)\, s_{jt}}\right], \quad i \neq k. \tag{3}$$

Table 1: Below-cost share and implied elasticity

ATC1 class	N drugs	Below-cost share	$\lvert\varepsilon^{\text{own}}\rvert$
H (systemic hormones)	6	0.45	1.16
L (antineoplastics)	2	0.40	0.31
S (sensory organs)	3	0.30	0.19
C (cardiovascular)	38	0.26	1.35
G (genito-urinary)	42	0.23	1.28
A (alimentary tract / metab.)	20	0.13	0.78
D (dermatologicals)	1	0.10	3.47
B (blood / blood-forming)	10	0.10	1.24
N (nervous system)	59	0.07	1.08
R (respiratory)	22	0.06	2.15
M (musculo-skeletal)	14	0.06	1.65
J (anti-infectives, systemic)	5	0.01	3.26

Note: Below-cost share is the fraction of chain-drug-week observations in the pre-ban window with $p_{ijt} < c_{jt}$; $\lvert\varepsilon^{\text{own}}\rvert$ is the class-mean of $\lvert\varepsilon^{\text{own}}_{ijj,t}\rvert$ from (2) at median pre-ban prices. Spearman rank correlation across the 12 classes: $\rho = -0.62$.

The pre-ban share-fit correlation 0.73 reported in main-text Table 7 is below the post-ban value of 0.86. This appendix decomposes the gap. The three relevant sub-regimes are: pre-ban competitive weeks (panel-wide below-cost prevalence ≤ 25%), pre-ban war weeks (> 25% below-cost prevalence), and post-ban weeks. At the headline parameters $(\hat{\alpha}_0, \hat{\alpha}_1, \hat{\sigma}) = (0.103, -0.074, 0.393)$:

Table 2: Share-fit decomposition by regime

Regime	N	Corr($\hat{s}_{i|jt}$, $s_{i|jt}$)
pre-ban competitive	44,421	0.67
pre-ban war	18,564	**0.88**
pre-ban (all)	62,985	0.73
post-ban	41,769	0.86

Note: Pre-ban "war weeks" are pre-ban weeks with panel-wide below-cost prevalence above 25% (28 weeks). Pre-ban "competitive" are the remaining pre-ban weeks.

The pre-ban war weeks fit better than the post-ban regime ($\rho = 0.88$ vs. 0.86). The fit gap is therefore concentrated in the pre-ban competitive weeks, where the static nested-logit struggles to match the strategic dimension of loss-leader behaviour: in those weeks each chain was independently selecting which drugs to push below cost and by how much, drug-by-drug and week-by-week. These are coordinated intertemporal pricing decisions that do not enter the right-hand side of (2); the nested-logit sees them as residuals ξ_{ijt}. Post-ban, the three chains converge to a single coordinated tier and the strategic dimension flattens.

I tested allowing the price coefficient to differ in war weeks, $\alpha_t = \alpha_0 + \alpha_1 D_t + \alpha_2 W_t$ where W_t is the war-week indicator, by a grid search that maximises pre-ban share correlation. The optimum is $\hat{\alpha}_2 = +0.05$ but the gain in pre-ban ρ is negligible. Removing the regime split entirely and fitting a whole-period $(\hat{\alpha}, \hat{\sigma})$ by share-fit grid search yields $(\hat{\alpha}, \hat{\sigma}) = (0.05, 0.15)$ with whole-period $\rho = 0.80$, marginally above the pooled spec's $\rho = 0.78$. The modest gain confirms that the binding constraint is the within-competitive-week strategic dimension, not the between-regime price coefficient. A fuller resolution would require a dynamic model of consumer search or chain pricing choices; I leave that to future work and proceed with the single $(\hat{\alpha}_0, \hat{\alpha}_1, \hat{\sigma})$ for the welfare calculation in §5, which uses tier-level prices where shares are stable.

A.1 The product-firm demand the structural model inherits

The dynamic game inherits demand at the chain-by-drug level. Each chain i's quality intercept $\hat{\phi}_{ij}$ is recovered from the within-nest shares $\hat{s}^{w}_{i|j}$ observed in the coordinated window via (1) at the common coordinated price; the instrumented coefficient $\hat{\alpha}_j$ governs the price response and the market size N_j is anchored on a normal-period median quantity. Two fit checks support the mapping: out-of-sample market-share fit and quantity fit for the objects the game needs.

Per-drug specification. Holding the nesting parameter at the pooled $\hat{\sigma} = 0.393$, I estimate the per-drug $\hat{\alpha}_j$ for each of the 222 drugs by two-stage least squares on the competitive (cartel-excluded) daily panel ($\approx$ 3,288 obs per drug across three chains), instrumenting each drug's price (and its post-ban interaction) with the two rival-chain prices.[1] The within-transform removes chain (Firm) effects only: a two-way Firm×Day demeaning would force the three chains' prices to sum to zero within each day, making own price a mechanical function of the rivals and the rival instrument degenerate. The within-drug estimating equation is

$$\tilde{Y}_{ijt} \;=\; -\,\alpha_j\, P_{ijt} \;-\; \gamma_j\,(P_{ijt} \times \text{PostCartel}_t) \;+\; \phi_i \;+\; \varepsilon_{ijt}, \tag{4}$$

with $\tilde{Y}_{ijt} = Y_{ijt} - \hat{\sigma} \ln s_{i|jt}$, $\text{PostCartel}_t = 1$ for days 829–1,096, ϕ_i a chain fixed effect, and P_{ijt} (and its interaction) instrumented by the rival prices. Identification of α_j relies on within-drug, cross-chain price dispersion in the competitive windows. The ban interaction γ_j enters only as a control, so $\hat{\alpha}_j$ is the clean pre-ban within-window price slope; the post-ban $\hat{\alpha}_j$ is re-estimated with the Salcobrand-cost instrument, since post-ban rival prices are coordination-contaminated, and the structural model scales each $\hat{\alpha}_j$ to its post-ban value by the common pooled ratio $\alpha^{\text{post}}/\alpha_0 = 0.29$ (§5). The pooled analog of the ban shift is $\hat{\alpha}_1 = -0.074$.

Market share. Intercepts recovered at the coordinated price predict the within-nest shares in two regimes they were not fit to (the price war and the supra-competitive rent) to a median absolute error of 1.6 percentage points (mean 2.5; 90th percentile 5.6) across the 216 drugs. The coordinated window is in-sample and reproduces exactly.

Quantity. With N_j anchored on the normal-period quantity, the implied three-chain quantity $N_j \sum_i \hat{s}_{ij}$ matches the observed quantity to a median ratio of 0.97, 0.99, and 0.98 at the war, coordinated, and rent tiers (interquartile range $\approx [0.85, 1.13]$; log-correlation 0.75–0.85). Anchoring instead on the absolute-peak quantity (the price-war stocking-up spike) over-states quantity by a third and is rejected.

The chain-specific undercut gains reduce to the symmetric gain when the three chains' within-shares are equal and correlate 0.94 with it across drugs, so the product-firm layer nests the symmetric mapping as a special case.

[1] The nesting parameter σ governs the single chain nest (the three chains compete for the same drug), so it is a market-level, not a drug-level, object. I therefore identify it once on the pooled pre-ban data (Panel A of main-text Table 6) and condition the per-drug slopes on it; a separate per-drug σ_j is not identified, as it would need within-drug substitution variation that three chains over a short window cannot supply. Estimating a common σ with heterogeneous price coefficients α_j is the standard sequential nested-logit approach [Berry, 1994]. Table 10 (Panel B) re-solves the structural model at $\sigma \in \{0.30, 0.50\}$: the μ band and the war-to-coordination flip are stable, so the fixed-σ choice does not drive the results.

Figure 1: Out-of-sample market-share and quantity fit

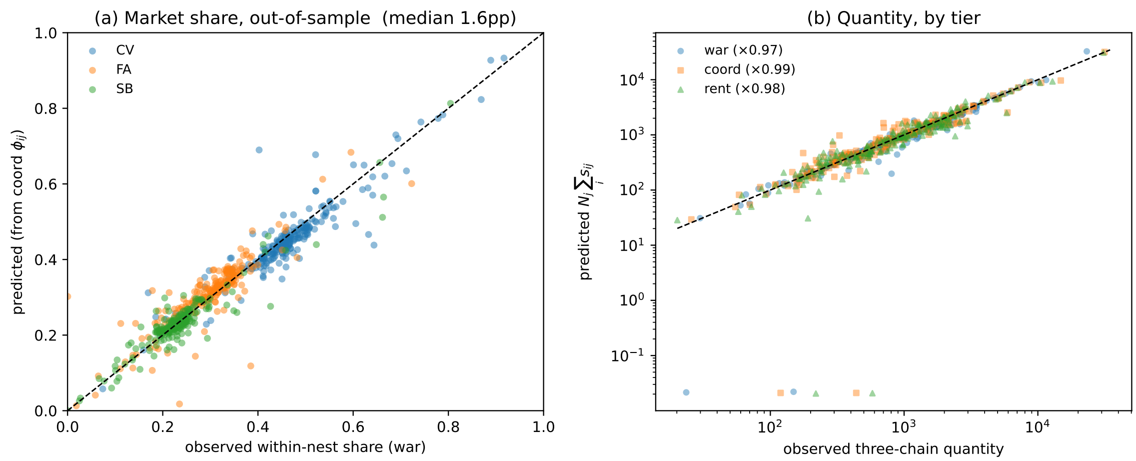

Note: Panel (a): predicted versus observed within-nest shares in the price war, using intercepts $\hat{\phi}_{ij}$ recovered from the coordinated window (out-of-sample); the 45° line is exact fit. Panel (b): predicted three-chain quantity $N_j \sum_i \hat{s}_{ij}$ versus observed, by tier (log scale). 216 drugs with a complete war / coordinated / rent price triple.

B The event ladder vs. the binary cartel indicator

B.1 Event coding versus the binary cartel indicator

The case data include a coordination indicator, $\text{col}_{fjt} \in \{0,1\}$, equal to one when chain f's price for drug j on day t is in the coordinated (post-increase) regime. This indicator records whether a firm has joined the coordinated price but not how far prices have risen: a $0 \to 1$ flip marks the first coordinated increase, yet the variable cannot distinguish a first increase from a second or third. The event ladder instead records the full sequence of coordinated increases for each drug: every transition from price tier ℓ to $\ell+1$, its date, and the firm that moved first.

Table 3 compares the two. The binary indicator coincides with most first-round onsets but with only 22% of the second- and third-round increases, and in 127 of the 137 multi-round drugs it records fewer steps than the event ladder. Because the mechanism I study is the sequential, multi-round escalation itself, the tier ladder, not a single on/off flag, is the economically relevant unit of observation. Figure 2 illustrates with four representative drugs: prices climb in discrete plateaus that the event ladder tracks step by step, while col collapses the entire episode into one block.

Event-time comparison. I further compare the event ladder with the FNE official cartel

Table 3: Event ladder versus binary indicator

	Value
Drugs in the coordination sample	222
Coded coordinated price increases	442
first-round onsets (tier $0 \rightarrow 1$)	293
later-round steps (tier $1 \rightarrow 2$, $2 \rightarrow 3$)	149
Drugs with ≥ 2 coordinated increases	137
col $0 \rightarrow 1$ flips (drug level)	249
First-round onsets matched by a col flip	188 (64%)
Later-round steps matched by a col flip	33 (22%)
Multi-round drugs where col records fewer steps	127 of 137
Leader agreement (col first-mover vs. coded leader)	57%

Notes: a col flip is matched to a coded event if the two occur within 15 days. Later-round steps have no counterpart in a binary indicator by construction.

Note: The event ladder vs. the binary col indicator

indicator $\text{col}_{ijt} \in \{0, 1\}$ at the drug level. For each drug j I define the FNE official event time as the first day on which $\text{col}_{\text{CV},jt} = \text{col}_{\text{FA},jt} = \text{col}_{\text{SB},jt} = 1$. Of the 205 drugs that have an event in both panels, the median difference (FNE day minus ladder day) is -2 days, with inter-quartile range $[-3, +5]$ days and Pearson correlation $r = 0.768$. Seventy-nine percent of drugs are within seven days of each other across the two panels, and 95% are within sixty days.

C Market-size robustness

The demand estimator of Section 4 uses $N_j^{\text{wk}} = 7 \max_t Q_{jt}^{\text{wk}} + 1$ in the weekly headline panel. The daily per-drug panel uses the corresponding weekly-equivalent normalization $N_j^{\text{day}} = 7 \max_t Q_{jt}^{\text{day}} / \hat{\rho} + 1$, with $\hat{\rho} = 0.92$ (three-chain market share, FNE 2008). I re-estimated the model under two alternative calibrations.

1. **Higher chain-share estimate.** Set $\hat{\rho} = 0.95$ (Díaz and Galetovic, 2015) in the daily per-drug normalization. This lowers the daily-panel N_j by 3.3%. Point estimates of $(\hat{\alpha}, \hat{\sigma})$ shift by less than ± 0.005.
2. **Constant national population.** Set $N_j = N^*/J$ for all j, with $N^* = 4 \times 10^6$ (Chile's adult prescription-buying population). $(\hat{\alpha}, \hat{\sigma})$ shift by less than ± 0.01.

Figure 2: Event taxonomy versus the binary indicator

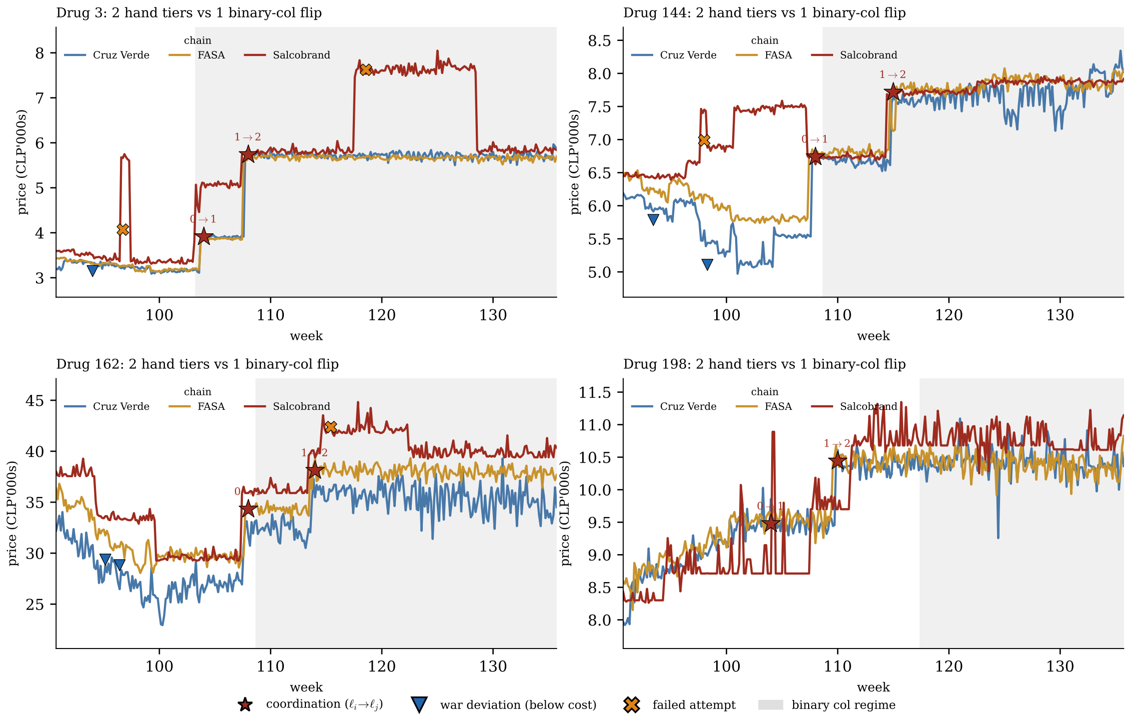

Note: Each event is marked at its own date and price by nature: coordination tier transitions (stars, labelled $\ell_i \to \ell_j$), below-cost war deviations (triangles), and failed attempts (crosses); post-coordination defections use diamonds where present. Prices move in discrete tiers that the event coding records individually, with their nature distinguished, whereas the binary indicator marks only the onset of the coordinated regime.

The own-price elasticity dispersion across drugs is stable across the three specifications, confirming that the chain-by-drug fixed effect ϕ_{ij} absorbs level mis-scaling of N_j.

D Data construction

The transaction data, obtained from the FNE under Chile's transparency law (Ley No. 20.285), cover all purchases of the 222 case products from the three chains over 2006–2008; each record carries the drug, chain, date and time, list price, transacted price, and units. I aggregate to weekly revenue-weighted averages: daily data are noisy at low volumes, while monthly aggregation would blur the 2–4-day coordination windows. Product attributes (active ingredient and ATC class; package size, pill count, brand-vs-generic status, prescription status) were hand-compiled and cross-checked from three public sources (an internal catalog, DrugBank, and the Chile-specific listing Farmazon) with the small share of conflicts ($< 3\%$) resolved in favour of the Farmazon listing. Wholesale costs are Salcobrand's wholesale prices for November 2007–May 2008, submitted to the Competition Tribunal as expert evidence [Tribunal de Defensa de la Libre Competencia (TDLC), 2012]. For chain-product-periods with no recorded sales I use the chain-product median price; periods with zero quantity are treated as missing, not as a zero price; and within-chain coordination statistics use only chains that stock the product.

Mapping to the 222-drug TDLC list. Of the 222 drugs named in the TDLC *Sentencia*, 220 (99.1%) show at least one successful three-chain coordination in the event panel; 211 clear the 15% coding threshold, and 9 are recovered only by the more permissive daily-panel detector or the later-window pass (small tier-0 → 1 events of 12–15%). The remaining two, Losopil (Mediapharm) and Progyluton (Bayer), show only single-firm unilateral lifts of 12–14%, with no event in which all three chains follow. Both count as coordinated in the official `col_ijt` indicator (a binary flip covering the full 222-drug list by construction), but my finer panel, which separates single-firm attempts from three-chain agreements, does not classify them as coordinated. The implied three-chain coordination rate on the TDLC list is $220/222 \approx 99.1\%$.

E The structural estimation sample

The dynamic model of §5 is estimated on the $J = 220$ drugs that complete the war-to-coordination price transition the model is built around. A drug enters the structural panel only if the event ledger records all three of (i) a war/deviation price (a pronounced dip during the price-war window), (ii) a Tier-1 coordinated price, and (iii) an imputed cost; the Tier-2 rent price is imputed when absent, and the per-drug price coefficient α_j falls back to its type median when the demand regression is uninformative.

An initial pass over the event ledger excluded fifteen drugs. Cross-checking each against its raw price path showed that nine of them do complete the transition, a marked dip during the war window followed by a sustained higher level, that the ledger had failed to record; I reclassified these (restoring the missing war and/or Tier-1 events), raising the sample from 207 to 216. A second pass then recovers four more (Table 5): Carvedilol, Chlorpheniramine Maleate, Ethinyl Estradiol, and Oxolamine Citrate are all coded three-chain coordinations (Figure 3 shows each stepping up to a sustained higher level in 2008) whose only missing ledger leg is a flagged war-deviation price. I impute that leg from the drug's demand-panel war-window median (at the 0.82 ratio that aligns the demand-window median with the recorded war-deviation price across the in-sample drugs; `full_structural.build_panel`), which brings the structural sample to $J = 220$, exactly the descriptive count of coordinated drugs in main-text Table 2. This price-path audit doubles as a completeness check on the event coding. Two case drugs remain excluded: the Neuractin presentation of Valproic Acid (#202) coordinates in early 2008 but then defects after a monotonic price rise, with no preceding loss-leader war to identify the transition; and Estradiol Valerate/Norgestrel (#215) never reaches a sustained Tier-1 level. Excluding these two is conservative for the coordination-timing moments.

Table 4 lists the moments the SMM criterion targets and the parameter each primarily identifies, the mapping summarised in §5.

Table 4: Targeted moments and what they identify

Group	Moment	Primarily identifies
Counts	first ($0 \to 1$) and second ($1 \to 2$) coordination totals	λ; rent penalty χ
Timing	onset week, peak week, wave SD	t_0, λ
	$1 \to 2$ lag behind $0 \to 1$	rent penalty χ
Laboratory	lab-wave SD; labs coordinated by wk. 104; within-lab SD	reach-out λ, τ_e
	corr(log lab size, coordination week)	reach-out λ, τ_e
War floor	below-cost war markup	μ (joint)
War series	weekly below-cost deviations: pre / mid / late / post	g
Micro	corr(pre-ban cheat freq., G_j^{01})	benefit-ordered war
	α_j/margin gap: rent-takers vs. not	rent selectivity
Belief	failed attempts by period (war high, cartel ≈ 0, post-investigation rebound)	war prior (a_0, b_0), signal ζ, skepticism m
Rent	rent steps by period (cartel + post-investigation deepening)	rent penalty χ, skepticism m
Leadership	leader share by period (SB wave, CV late rent)	outside option ψ (calibrated)

E.1 Bootstrap standard errors and the identification of the belief jump

The step-function criterion has no moment Jacobian, so the standard errors in main-text Table 8 are a nonparametric bootstrap. I resample the 220 drugs with replacement, recompute the empirical moments, and re-estimate the model on each of 300 replications, each warm-started at the headline estimate. The standard error is the standard deviation across the draws.

The ad-ban jump ζ is the one parameter a symmetric standard error does not summarize, and the reason is informative. ζ is identified only from below. No resample places it under about 180, so the data demand a large focal jump but cannot bound it from above.

Table 5: Borderline case drugs: recovered or excluded

Drug ID	Active ingredient	Type	Status	Ledger note
169	Carvedilol	Chronic OTC	Recovered	war price imputed
182	Chlorpheniramine Maleate	Acute OTC	Recovered	all legs imputed
206	Ethinyl Estradiol	Acute Rx	Recovered	war price imputed
210	Oxolamine Citrate	Acute OTC	Recovered	war price imputed
202	Valproic Acid (Neuractin)	Chronic Rx	Excluded	coordinates, then defects; no war
215	Estradiol Valerate/Norgestrel	Chronic Rx	Excluded	never reaches Tier-1

Note: Four of the six are recovered into the structural sample, each a coded three-chain coordination whose only missing ledger leg is a war-deviation price imputed from the demand panel. This raises J from 216 to 220, the full count of coordinated drugs; two remain excluded. "Type" is the chronic/acute × prescription/OTC classification.

Figure 3: Six borderline drugs: four recovered, two excluded

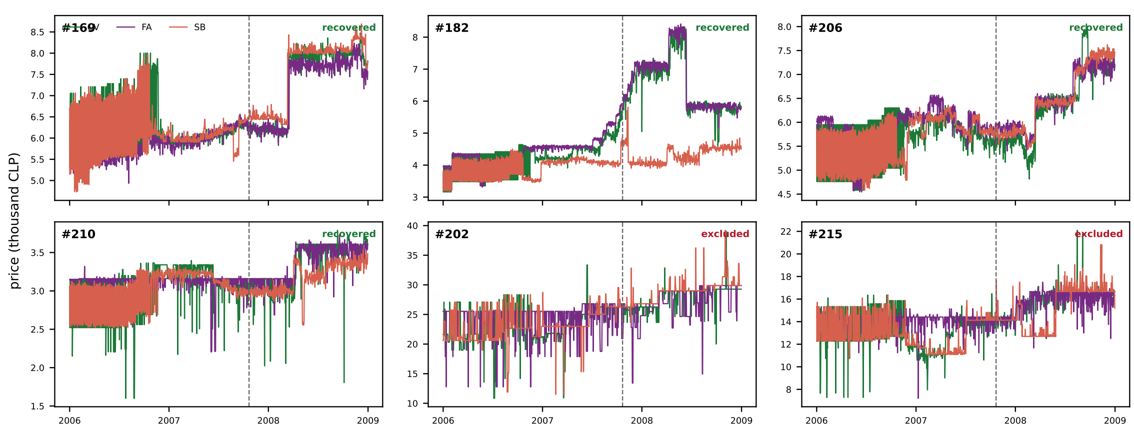

Note: All six borderline drugs I adjudicated are plotted: the four recovered (green) I admit to the sample and the two excluded (red) I drop, so only the two red drugs leave the sample (the figure shows both decisions, not only the drops). Each panel plots the three chains' daily price (thousand CLP, cleaned of abnormal jumps as in the demand panel) for one drug; the dashed line marks the advertising ban. The four recovered (green) each step up to a sustained higher level in 2008 (a coded three-chain coordination), but their war-period dip is shallow or unrecorded, so I impute the war-deviation price from the demand panel and admit them, taking the structural sample to J = 220. The two labelled excluded (red) stay out: #202 rises monotonically and then defects without a preceding loss-leader war, and #215 never reaches a sustained Tier-1 level.

Table 6 shows why. With no jump the coordination wave collapses, but from $\zeta = 150$ to $\zeta = 2000$ the wave saturates at about 215 first-coordinations, because once the jump is large enough to flip beliefs the posterior saturates and larger jumps are observationally equivalent. The bootstrap distribution of ζ is right-skewed with a sharp floor (median 315, interquartile range [296, 370], no draw below 179) and a long upper tail, so I report it as a lower bound, $\zeta \gtrsim 180$.

Table 6: The coordination wave saturates once the jump is large

ζ	SMM loss	First coordinations	Rent steps
0	2389	0	0
150	142	213	123
307	**16**	**215**	**141**
500	28	215	146
800	43	215	149
2000	68	216	153

Note: The headline model re-simulated with only ζ varied, all other parameters held at their estimates. With no jump ($\zeta = 0$) coordination collapses. From $\zeta = 150$ to $\zeta = 2000$ the wave is saturated, so the data fix ζ sharply from below and only loosely from above. The estimate $\zeta = 307$ minimizes the loss.

F Robustness to the demand specification

The core mechanism is robust to how the demand side is specified, but the exercise also shows why the headline keeps per-drug demand heterogeneity. Table 7 re-estimates the model under three demand specifications on the identical headline criterion: the heterogeneous per-drug $\hat{\alpha}_j$ from the rival-price IV (Het. IV, the headline; main-text Table 6); the same heterogeneous $\hat{\alpha}_j$ from the OLS (Het. OLS); and a homogeneous demand that gives every drug the median $\hat{\alpha}$. The store-traffic level ($\bar{\mu} \approx 9.0$–9.3) and the leadership ordering are stable across all three: Salcobrand leads the cartel wave (SB-share 0.68–0.69) and Cruz Verde takes over the post-investigation rent (CV-share 0.67–0.69). The two heterogeneous specifications give the same structural estimate and fit comparably (loss 16.6 and 17.3; the two $\hat{\alpha}_j$ series are nearly collinear). Homogeneous demand fits worse (loss 27.4): without the per-drug heterogeneity the rent over-deepens and arrives too late (169 steps against 143, a $1 \rightarrow 2$ lag of 13.3 against the data's 8.0 weeks), so the per-drug heterogeneity is what sizes and times the rent.

Table 7: Structural estimates across three demand specifications

	Data	Het. IV	Het. OLS	Homog.
Panel A. Estimated parameters (SMM)				
$\bar{\mu}$ mean store-traffic value		9.26	9.26	9.02
λ reach-out scale		0.63	0.63	0.34
τ_e reach-out temperature		124	124	54
t_0 organization onset (wk)		103	103	102
a_0, b_0 belief prior		1.80, 129	1.80, 129	2.11, 129
ζ ad-ban pseudo-successes		307	307	325
χ rent-tier discount		2.08	2.08	2.59
m enforcement skepticism		14.8	14.8	17.1
g Aug. campaign strength		0.57	0.57	0.64
SMM loss		**16.6**	**17.3**	**27.4**
Panel B. Fit on selected moments				
First coordinations	220	214	220	220
Rent steps	143	141	140	**169**
Failed war attempts	123	110	110	126
Cartel successes	284	289	292	304
Post-late successes	49	34	34	54
Peak week	107	107	107	106
Wave SD (weeks)	4.6	4.5	4.6	4.0
$1 \to 2$ lag (weeks)	8.0	8.1	8.1	**13.3**
Between-lab var. share	0.59	0.57	0.61	0.48
Spearman(size, week)	−0.30	−0.24	−0.28	−0.24
SB leads cartel wave	0.74	0.69	0.69	0.68
CV leads post-invest. rent	0.65	0.68	0.67	0.69

Note: each column re-estimates the eleven structural parameters (symmetric $\bar{\mu}$ plus the belief, timing and rent parameters) by SMM on the headline criterion; the engagement window $w = 26$, δ, τ and the holiday factor are fixed. Het. IV and Het. OLS differ only in the per-drug $\hat{\alpha}_j$ (main-text Table 6, Panel B: rival-price 2SLS vs. OLS); the two series are nearly collinear, so they yield the same structural estimate. **Homog.** sets every $\hat{\alpha}_j$ to the median and selects rent-eligible drugs by the flow condition alone. The main text reports the rival-IV estimate; the mechanism (belief transition, lab-organized wave, and the Salcobrand→Cruz Verde leadership handoff) survives all three demands, only the homogeneous spec mistiming and over-deepening the rent.

Estimating the collapse. The headline applies the pooled post-ban collapse $\alpha^{\text{post}}/\alpha_0 = 0.29$ (main-text Table 6) to each drug's pre-ban $\hat{\alpha}_j$. As a further check I estimate the collapse as a single parameter jointly with the per-drug pre slopes, in two forms: a proportional reduction $\alpha_j^{\text{post}} = \kappa\,\hat{\alpha}_j$ ($\hat{\kappa} = 0.347$, SE 0.006) and an additive reduction $\alpha_j^{\text{post}} = \hat{\alpha}_j - \delta$

($\hat{\delta}$ = 0.053, SE 0.001). Re-estimating the structural model under each gives SMM loss 16.6 (proportional, matching the headline) and 25.2 (additive). The proportional collapse is virtually the headline demand. The additive collapse fits worse: it misses part of the first-coordination wave (194 against 220), although it still matches the rent count closely (146 against 143) and keeps the Salcobrand-first cartel-wave leadership (0.68). Thus this check supports the mechanism, but the proportional collapse is the closer empirical description.

G Structural robustness checks

Each check re-simulates the headline at the estimate with one feature switched off, reporting the fit moments of main-text Table 8 (Panel B). Table 8 collects the three value-function checks that justify the dynamic structure; the belief and payoff knockouts are in main-text Table 9.

Myopic decisions ($\delta \to 0$). A myopic firm weighs only the current period, so the continuation value ΔV^1 drops out of the incentive constraint. Coordination then survives only where holding is already a one-shot best response: the count falls to 151 of 220 and the rent to 107 of 143, and because the escape value no longer differs across chains the leadership collapses toward chance (Salcobrand 0.26). Escaping the war and deepening the rent each trade a current cost for a future payoff, so the forward-looking value is necessary; the fit degrades to a loss of 555.

Perfect foresight of the ban. The headline treats the November injunction as an unanticipated public event. The alternative lets firms anticipate it from the first public signal, the 7 September CONAR ruling, so the coordination belief jumps there. Coordination then starts in September, two months before the binding event (onset week 88 against the data's 102), and because the pre-ban demand still rewards undercutting the premature attempts appear as failures (war-period failed lifts 717 against 123 in the data). The data place the operative shift at the binding November event, not the foreseeable precursor; the fit degrades to a loss of 2550.

Absorbing value functions. The headline values each state by forward-simulating the model's own transitions, so a firm in the war anticipates it may coordinate and a coordinated firm anticipates the rent may revert. Valuing each tier instead as a permanent perpetuity ($V^x = \Pi^x + \delta V^x$), as if a coordinated drug never reverts and a rent never fails, overstates the tier values and mistimes the rent (rent steps 138 of 143, loss 23). This is the mildest of the three, as it should be: it refines the continuation value rather than removing a channel, but the transition-aware forward simulation is what reproduces the post-investigation dynamics.

Leadership. Two further checks confirm that leadership is the calibrated outside option, not the spillover. Setting ψ symmetric collapses Salcobrand's lead to about 27%

while leaving the coordination count untouched, the signature of a leadership channel decoupled from the incentive constraint. Restoring the asymmetry through a chain-specific store-traffic spread μ_i instead reproduces the Salcobrand-first split only by fitting μ_i to the leadership it then explains, the circularity a calibrated outside option avoids (§5).

Table 8: Value-function robustness checks: fit moments

Moment	Data	Estimated	Myopic ($\delta \to 0$)	Perfect foresight	Absorbing
First coordinations ($0 \to 1$)	220	214	151	188	214
Rent steps ($1 \to 2$)	143	141	107	108	138
Onset / peak week	102/107	103/107	105/118	**88**/90	103/105
$1 \to 2$ lag (weeks)	8.0	8.1	23.7	28.8	7.5
Cartel successes	284	289	95	81	298
Failed lifts: war	123	110	110	**717**	111
SB leads the cartel wave	0.74	0.69	**0.26**	**0.34**	0.68
SMM loss	—	**16.6**	**555**	**2550**	**23**

Note: each column re-simulates the headline at the estimate ($B = 240$) with one feature of the value function switched off. Myopic: $\delta \to 0$ (no continuation value). Perfect foresight: the coordination belief jumps at the 7 September CONAR precursor rather than the binding November ban. Absorbing: each tier valued as a permanent perpetuity. Compare "Estimated" to main-text Table 8, Panel B.

The cost of leading. A chain that raises a drug first is the lone high-price seller, and a rival that stays low steals a share Δs_j of the drug's volume; that share loss is the cost of leading. Figure 4 plots its distribution across the 220 drugs, before and after the ban. While comparative-price advertising is on the cost is large (median $\Delta s \approx 0.14$), so the lone mover is undercut and the war persists; the ban collapses it (median ≈ 0.04), so leading becomes cheap. The cost never gates coordination, however. The discounted escape value exceeds even the pre-ban cost by an order of magnitude (§5, the slack incentive constraint), so the most-exposed chain leads despite it, and the coordination count is the same whether the store-traffic spillover is large, small, or shut off entirely (main-text Table 9).

Figure 4: The cost of leading

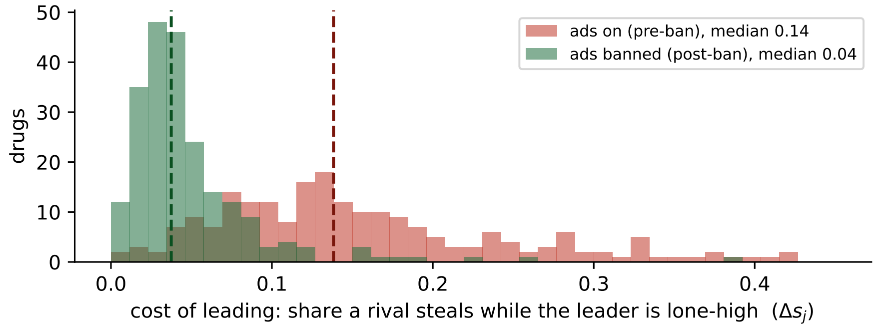

Note: distribution across the 220 drugs of the cost of leading, the within-period share Δs_j a rival steals while the leader is the lone high-price seller, computed from the estimated demand. Dashed lines mark the medians. The ban collapses the cost from a median 0.14 (ads on) to 0.04 (ads banned).

The cost of leading is overwhelmingly a loss-leader cost. Table 9 splits the median cost into the drug margin lost on the customers who defect to the cheaper rivals and the store traffic those customers take with them. The store-traffic piece is about eighty-five percent of the total, so a chain that leads gives up little drug profit and mostly forgoes the basket sales of the shoppers it loses. The ban collapses both pieces with the share loss. The net cost, after the higher margin a chain earns on the customers it keeps by leaving the below-cost war, is smaller still and turns negative once the ban lands, so leading is then profitable on its own.

Table 9: The cost of leading, in money

Per drug customer, per week (median)	Ads on	Ads banned
Share lost by leading	0.11	0.04
Lost drug margin on defectors (CLP)	185	67
Loss-leader, store traffic μ (CLP)	1,050	395
Gross cost of leading (CLP)	**1,235**	**462**
Net cost, after retained-margin gain (CLP)	405	−305

Note: Median across the 220 drugs of the within-period cost of leading a coordinated increase, decomposed from the estimated demand. The store-traffic value sets the per-customer scale ($\hat{\mu} \approx$ 9,300 CLP). "Gross" is the value of the defecting customers, their drug margin plus the store traffic they bring; "net" subtracts the higher margin the leader earns on the customers it retains by exiting the below-cost war. The figure reports the unconditional share-loss median; this table reports the median after the money decomposition, so the ads-on share row need not equal the figure's 0.14 line. Both are an order of magnitude below the coordination prize $\delta\Delta V^1 \approx$ 11,100 CLP, so the incentive constraint stays slack.

Robustness of the coordination flip. Table 10 reports the threshold cushion behind the parameter-free flip discussed in §5. Panel A varies the size of the post-ban drop in α; Panel B perturbs the wholesale-cost and nesting calibration. The band $[\mu^*_{\text{pre}}, \mu^*_{\text{post}}]$ never closes.

Table 10: Robustness of the coordination flip

Panel A. Cushion vs. the size of the α drop			
post-ban α	μ^*_{post}	cushion $\mu^*_{\text{post}}/\mu^*_{\text{pre}}$	band
0.029 (cartel-excl. est.)	30.3	3.5×	open
0.060	14.6	1.7×	open
0.100 (rival-IV upper bound)	8.8	1.0×	open (sliver)
Panel B. Sensitivity of $[\mu^*_{\text{pre}}, \mu^*_{\text{post}}]$ to the calibration			
perturbation	μ^*_{pre}	μ^*_{post}	ratio
baseline	8.6	30.3	3.5×
wholesale cost −10%	7.9	29.6	3.8×
wholesale cost +10%	9.3	31.0	3.3×
nesting σ = 0.30	9.8	34.8	3.6×
nesting σ = 0.50	7.2	25.0	3.5×

Note: $\mu^*(\alpha)$ is the store-traffic value at which undercutting the cartel ties coordinating; the flip requires $\mu^*_{\text{pre}} < \mu < \mu^*_{\text{post}}$. Panel A varies the size of the post-ban drop in α ($\mu^*_{\text{pre}} = 8.6$ fixed at $\alpha = 0.103$); Panel B perturbs the calibration.

Belief and payoff specifications. The main-text Table 9 re-estimates the model with each belief channel and each payoff ingredient switched off in turn, the evidence behind §5 and the per-moment companion to main-text Figure 7. The focal jump dominates: dropping the Bayesian learning (jump-only) still reproduces the cartel wave and fits essentially as well as the headline (loss 21 against 16.6), so the event-tally learning is a minor refinement. The focal jump itself is what starts the wave: with it switched off ($\zeta = 0$, learning-only) the coordination never forms in-sample, a degenerate corner (0 coordinations, loss ~ 2300). The commonly observed ban, not the slow tally, is what selects the coordinated equilibrium. The outside option's separate role (selecting who leads) is the re-simulated check in Table 8 (RC3: without it leadership collapses to uniform).

The rent's belief gate and pre-ban feasibility. The rent push (15) deepens an eligible coordinated drug ($\Delta V_j^2 > 0$) at a discounted belief gate $b_t^{\text{rent}} = e^{-\chi} b_t^{\text{coord}}$ times the learned enforcement belief r_t. Table 11 probes both. The belief gate is essential and times the rent: forcing $b_t^{\text{rent}} \equiv 1$, the deepening starts eight weeks early and produces too many (184 steps against 141), and the loss rises more than fivefold (88 against 16.6). The rent's pre-ban feasibility, by contrast, is not set by the coordination belief: forcing $b^{\text{coord}} \equiv 1$ over-produces the $0 \to 1$ wave before the ban (215 coordinations, against zero in both the data and the estimate) yet leaves the rent at zero before the ban, exactly as the estimate has it. The rent only deepens once the demand collapse makes the further increase pay; the post-ban flow eligibility, not a separate static-best-response gate, screens which drugs deepen.

Table 11: The rent belief gate and pre-ban feasibility

Moment	Data	Estimated	$b^{\text{rent}} \equiv 1$	$b^{\text{coord}} \equiv 1$
$0 \to 1$ before the ban	0	0	0	**215**
Rent steps before the ban	0	0	0	0
Rent steps ($1 \to 2$)	143	141	**184**	145
Rent peak week	117	115	**107**	—
SMM criterion (loss)	—	**16.6**	**88**	**4154**

Note: $b^{\text{rent}} \equiv 1$ drops the rent's discounted belief gate (re-estimated under the identical criterion); the rent then over-produces and starts early, so the gate is essential and times the rent. $b^{\text{coord}} \equiv 1$ forces the coordination belief to one (the no-belief column of main-text Table 9): the $0 \to 1$ wave over-produces before the ban, but the rent stays at zero before the ban because demand-side rent eligibility remains false before the ban. The rent still uses the discounted rent-belief gate for timing, but it carries no separate static-best-response gate; the eligibility already screens drugs that would statically defect.

G.1 Macro and ownership controls

Two contemporaneous shocks could in principle confound the structural reading: the August 2007 acquisition of Salcobrand by Empresas Juan Yarur and the associated management change, and the Chilean inflation of 2007–2008, which neared 10% at its mid-2008 peak. Table 12 re-estimates the model against each.

Ownership. The headline makes Salcobrand lead the cartel formation through a fixed outside option $\psi_{SB} = 0$: relative to FASA's international operations and Cruz Verde's Socofar structure, it is the most exposed domestic pharmacy chain in the calibration. The August 2007 acquisition is the institutional counterpart of that calibration: a leveraged buyer inherited an intensifying below-cost war, and Salcobrand led the December escape four months later. The timing supports the channel without making the acquisition the timing node: the cartel forms at the ban, not at the acquisition, so the buyout does not create the cartel; it rationalizes why Salcobrand, not Cruz Verde, is the eager leader. Freeing ψ_{SB} and re-estimating returns it to 0.00, with the fit and the leadership ordering stable (loss 16.6): the boundary re-estimation supports the calibrated corner.

Inflation. The deviation gain that drives the model, $G_j^{01} = \mu\,\Delta s_j + \Delta(\text{margin})$, runs on price–cost margins and within-market relative prices (the rival-price instrument differences out common shocks), so inflation that lifts prices and costs together is, to first order, neutral. The magnitudes confirm it: the war-trough-to-coordinated jump is about 45% over the four-month formation window, against roughly 4% inflation over those months. Re-estimating on prices deflated to the war base by the Chilean CPI (the coordinated tier by 4%, the rent tier by 10%) leaves the war below cost (markup −0.083 against −0.091) and the coordination count, the rent, and the leadership intact, at a slightly higher loss (17.7): the cartel is a real markup, not an artifact of the 2007–2008 inflation.

Table 12: Macro and ownership controls

Moment	Data	Headline	Real prices	Free ψ_{SB}
$\bar{\mu}$ store-traffic value		9.26	9.22	9.26
ψ_{SB} SB outside option		0 (fix)	0 (fix)	**0.00**
War markup $(p-c)$	−0.10	−0.091	−0.083	−0.091
First coordinations	220	214	216	214
Rent steps $(1\to 2)$	143	141	141	141
SB leads cartel wave	0.74	0.69	0.66	0.69
CV leads late rent	0.65	0.68	0.68	0.68
SMM criterion (loss)	—	**16.6**	**17.7**	**16.6**

Note: Real prices deflates the coordinated and rent tiers to the war-price base by the 2007–2008 Chilean CPI (∼0.7%/month: −4% on the coordinated tier, −10% on the rent) and re-estimates. Free ψ_{SB} frees Salcobrand's outside option (fixed at 0 in the headline) as a twelfth parameter. Both leave the war below cost, the coordination count and timing, and the Salcobrand→Cruz Verde leadership handoff stable.

G.2 The war forms endogenously, and experimentation cannot escape it

The main-text Table 9 shows the focal jump is essential and the Bayesian learning a contributing refinement. To see why, I roll the belief and price dynamics from January 2006 (a year before the structural window) through the war and into the cartel (Figure 5). The war is the low-belief one of two coexisting equilibria, and it forms on its own. High-belief coordination is dynamically sustainable throughout (§5), so the question is never whether the chains can coordinate but which equilibrium is selected. Through 2006 the war is fought by below-cost discounting: a chain cuts to steal store traffic, the rivals match within days, and the cut springs back (Fact 1). This is the war, and no chain yet holds a coordinated increase. As the below-cost pain deepens into 2007, clean attempts to lift out of it appear; but pre-ban every lift is undercut ($G_j^{01} > 0$) and reverts, so none sticks. The simulation reproduces the record on both counts: about 86 below-cost reversions in 2006 and 125 failed lifts in 2007, against 87 and 123, with zero successes. Discount or failed lift, every pre-ban reversion teaches the same lesson, that a raise will not be matched, so the belief never climbs above its low prior: it drifts from about 0.014 to about 0.005 by the ban, and the pessimism is self-fulfilling. The 6 November ban then delivers the focal jump and flips $G_j^{01} < 0$; the belief jumps sharply and the cartel forms.

What the focal jump adds: formation on the observed timescale. Suppose the ban's payoff effect occurs but there is no focal jump ($\zeta = 0$), so the chains raise the belief through the tally alone. Post-ban the static flip ($G_j^{01} < 0$) makes a coordinated raise a best response for the static-Nash majority, and a co-led raise faces only the b_t^{3-k} gate for k simultaneous leaders, far weaker than the single-leader b_t^2, so coordination is not mechanically ruled out. But it is far too slow to explain the record. Re-estimated without the jump, the wave never forms within the sample: no coordinated wave forms by the window's end, fitting an order of magnitude worse ($\sim$2300 against the headline 16.6). By pure experimentation the static-Nash majority would coordinate only over a horizon far longer than the observed window; its exact length turns on the learning rate, which the data identify only loosely, but on any reading it is far beyond the four-month wave, and the most-elastic drugs, whose G_j^{01} stays positive even post-ban, never coordinate without the belief at all. The focal jump's role is therefore to select the observed wave on the observed timescale: the commonly observed event resolves the higher-order belief in one step, compressing what bare experimentation would reach only much later into the sharp December onset the record shows.

The low belief is externally corroborated. The pessimism is not fit to the estimation window. Through 2006 the chains fought a deepening below-cost war and never sustained a coordinated increase, an observed coordination frequency of 0.00 against the estimated prior $a_0/(a_0 + b_0) = 0.013$; folding the 2006 record into the tally only lowers it. By the time the structural window opens, no raise had ever been matched.

Figure 5: The war forms endogenously from the belief dynamics

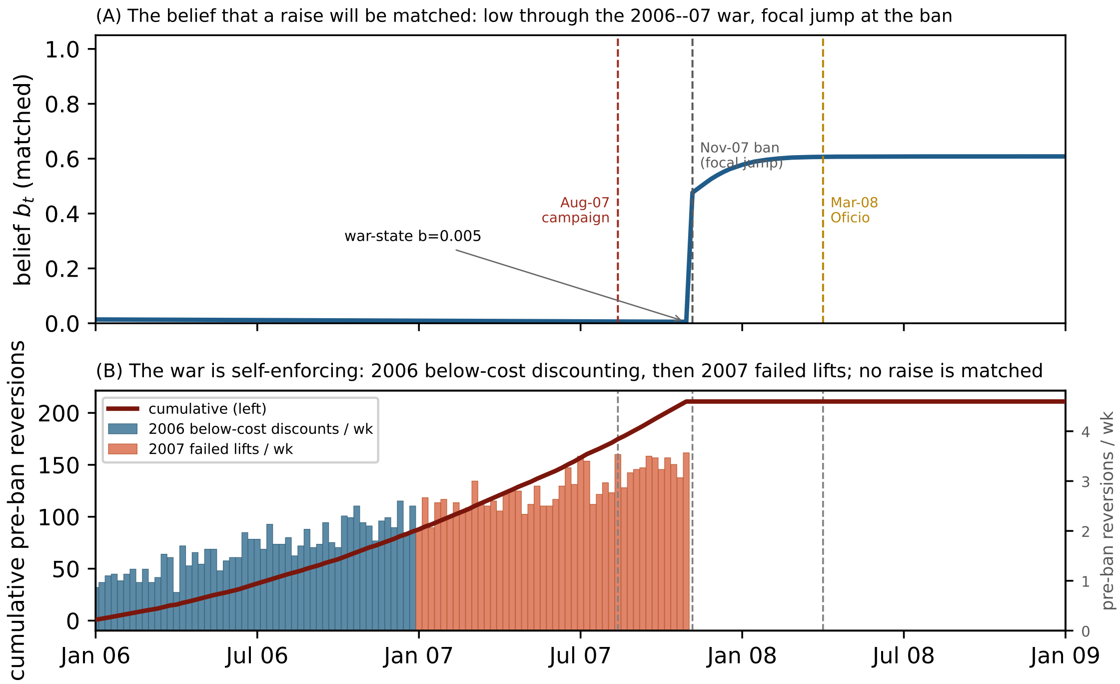

Note: belief and price dynamics rolled from January 2006 (week 0). **(A)** the belief that a raise will be matched: low and self-fulfilling through the 2006–07 war, then the focal jump at the 6 November ban. **(B)** the pre-ban price reversions that keep it low: 2006 is below-cost discounting (the war itself, Fact 1), 2007 the clean failed coordination lifts; the simulation matches the record (86/125 versus 87/123) with zero pre-ban successes. The reversion rate is calibrated to the two annual totals; the belief, hold and jump are the estimated model.

G.3 Leadership-dispersion robustness

The two leader-softmax dispersions τ_{lead} and τ_{rent} are calibrated to the leadership shares, so the share magnitudes are calibration targets, not predictions; the escape *ordering* is the prediction of the loss cushion ψ. Table 13 separates the two. Fixing $\tau_{\text{lead}} = \tau_{\text{rent}}$ at a common value, Salcobrand still leads the cartel wave throughout the tested range (69–80%), so the Salcobrand-first escape ordering is stable over that calibration range. The rent leader is more sensitive: Cruz Verde leads the held rent at the calibrated and low common scales, but the rent leadership flattens toward uniform at a high common dispersion. The robust object is the escape ordering; the share magnitudes and the rent leader are the calibration-dependent parts.

Table 13: Leadership robustness to a common softmax dispersion

Dispersion	SB leads cartel wave	CV leads the rent
τ_{lead}=3.0, τ_{rent}=0.12 (headline)	69%	68%
common τ = 0.12	73%	68%
common τ = 1.0	79%	37%
common τ = 3.0	76%	uniform (35/32/33)

Note: Leadership shares at the headline estimate, re-simulated with the two leader-softmax dispersions set equal. Salcobrand leads the cartel wave under every dispersion (the ψ ordering); the rent leader and the magnitudes depend on the calibrated scales.

G.4 Defection-rate robustness

The per-week down-cut rate λ_{cut} is a free scalar, but the post-ban cartel stability is an outcome of the slack gate, not of the rate. Table 14 perturbs λ_{cut} from half to double its estimate. The coordination count is stable over this range (first coordinations 214–215 of 220; cartel successes 286–292), while only the pre-ban war-deviation intensity scales with the rate. Post-ban the slack remains positive in these simulations, so undercutting does not pay over the tested rates; the rate governs how violent the *war* is, not whether the *cartel* survives.

Table 14: Cartel stability is stable over tested down-cut rates

λ_{cut}	First coord. (data 220)	Cartel succ. (data 284)	War deviations
0.044 (−50%)	215	286	253
0.088 (headline)	215	291	513
0.133 (+50%)	214	291	770
0.177 (+100%)	214	292	1019

Note: Headline estimate re-simulated with λ_{cut} perturbed. The coordination count barely moves; only the war's below-cost deviation count scales with the rate.

G.5 Dropping the noisiest drugs

The pooled demand collapse is used because the standalone per-drug post-ban coefficient is too thinly identified for some drugs: estimated drug by drug, 48 of the 220 receive a spuriously weak collapse. As a direct check that these noisy drugs do not drive the structural estimate, I drop all 48 and re-estimate on the remaining 172. Table 15 shows that the key estimates are stable. The store-traffic value is close to headline ($\bar{\mu}$ = 9.39 against

9.26), the down-cut rate is similar, and leadership shares move modestly. The model also reproduces the smaller-sample path in ratios: first coordinations 168/172, rent steps 114/117, and cartel successes 226/229. The criterion level is not comparable across the two samples because the count moments scale with the number of drugs, but the estimated primitives and the fit ratios are stable, so the weak-collapse drugs neither drive nor distort the estimate.

Table 15: Estimate is stable to dropping the 48 weak-collapse drugs

	Headline (J = 220)	Drop weak-collapse (J = 172)
$\bar{\mu}$ store-traffic value	9.26	9.39
λ_{cut} down-cut rate	0.088	0.090
SB leads cartel wave	69%	66%
CV leads post-rent	68%	66%
First coordinations (model/total)	214/220	168/172
Rent steps (model/data)	141/143	114/117
Cartel successes (model/data)	289/284	226/229

Note: The 48 drugs whose standalone per-drug post-ban collapse is weakest are dropped and the model re-estimated on the remaining 172. The store-traffic value and the leadership shares remain close; the count moments scale with the smaller sample.

H Example war attempts

Figure 6 shows six representative war attempts from the pre-ban period: episodes in which one chain raised a drug's price toward a coordinated level, the two rivals did not follow but undercut it, and the lift reverted to the below-cost war floor within weeks. These are the ledger's clean `failed_attempt` events (101 lift-then-revert episodes in the January–November 2007 war window, after setting aside 34 lower-amplitude price wobbles that do not clear a rise-and-revert threshold relative to the drug's own volatility), which the structural model reproduces through its belief-driven, sporadic war lifts (§5; 100 simulated). Salcobrand is the lifter in four of the six (Cruz Verde in the other two), consistent with its role as the coordination leader.

I Regulatory timeline and case outcome

The episode's documented chronology: a comparative-price advertising campaign by Cruz Verde (August 2007); the CONAR self-regulatory ruling (CONAR 704/07, 7 September 2007) and the binding Civil Court *medida precautoria* prohibiting comparative price

Figure 6: Six example war attempts

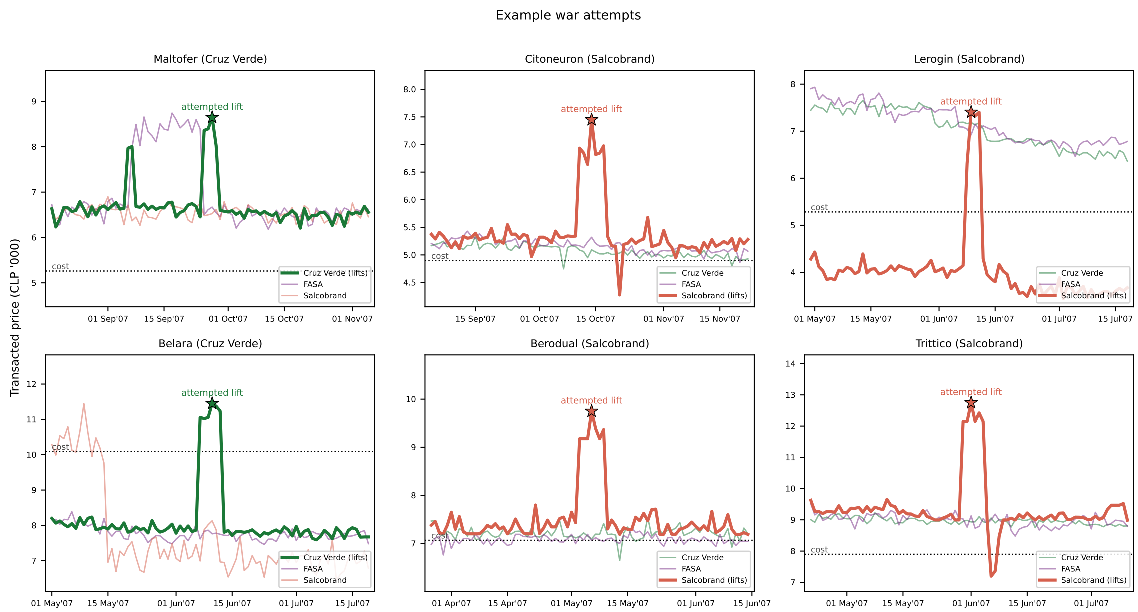

Note: In each panel the heavy line is the chain that raised price (marked ⋆ at the lift); the two light lines are the rivals, which stay near or below cost (dotted line) and undercut the lift, so it reverts within weeks. Transacted prices, three chains, ±6-week window around the event.

advertising (6 November 2007); the first three-chain coordinated increase (3 December 2007) and the wave running through April 2008; the FNE's complaint (Requerimiento, Rol C No. 184-08, 2008; Fiscalía Nacional Económica (FNE), 2008); and the Competition Tribunal's conviction on 206 drugs with fines of about US$38M (Sentencia No. 119/2012; Tribunal de Defensa de la Libre Competencia (TDLC), 2012). These documents provide the legal chronology used in the paper; the welfare accounting in §5 is computed from the transaction panel and the estimated demand system.

References

Steven T Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262, 1994.

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.

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`.

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`. Legal complaint, Tribunal de Defensa de la Libre Competencia.

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`.