Entry-Facilitating Merger Remedies: Theory and Evidence from the FCA–PSA Merger in Chile*

Jose Manuel Paz y Miño†

November 12, 2026

Abstract

We ask whether merger remedies that facilitate a competitor's entry protect consumers. When the merging parties propose a cap on an entrant's sales and approval precedes its entry decision, the remedy fails whenever the parties can secure approval with a cap that deters entry, even when entry could offset the merger's harm. Estimating demand and supply for the FCA–PSA merger in Chile, we find that the designated entrant, Toyota, could have offset the harm but stayed out, leaving consumers 5.51% worse off than absent the merger. No cap could have attracted Toyota while keeping the merger profitable in Chile.

*I am grateful to Ricardo Santolaya for his contributions to this project. I thank Andrés Gómez-Lobo and Juan Price for helping me get access to the data. I also thank Pierre Dubois, Eduardo Engel, Andres Gonzalez-Lira, Sebastian Fleitas, Vicente Lagos, Volker Nocke, Marcelo Olivares, Martin Peitz, Nicolas Schutz, André Stenzel and Matt Weinberg as well as other participants at Chile IO Day 2025 and EARIE 2026 for invaluable comments. I gratefully acknowledge financial support from the Chilean Government under the project Fondecyt de Iniciación en Investigación 2025 No. 11251217. This paper uses the public version of the Fiscalía Nacional Económica's decision and publicly available data. All errors are mine.

†Department of Economics, Universidad de Chile. Email: jpazm@fen.uchile.cl

1 Introduction

When a competition authority concludes that a merger raises competitive concerns, it can either block the transaction or approve it subject to conditions, known as remedies. Remedies are usually classified as behavioral or structural. Behavioral remedies regulate the future conduct of the merged firm, while structural remedies affect the very structure of the market. Agencies have consistently preferred structural remedies, particularly divestitures.[1] However, in many cases such remedies may not be feasible or proportionate. This typically occurs when the relevant assets lie outside the affected jurisdiction or serve a much broader market. Behavioral remedies alone are often insufficient to address the competitive concerns in these situations.

Faced with these limitations, authorities may rely on remedies that facilitate the entry or expansion of competitors through licensing, access or supply agreements, rather than through the transfer of assets.[2] Within this broader set of intermediate remedies, we focus on those specifically intended to facilitate the entry of new competitors (henceforth, "entry-facilitating" remedies).

Entry-facilitating remedies share a distinctive institutional structure. In many jurisdictions, the merging parties themselves offer the remedy,[3] and therefore set the terms of access, such as its price, its duration or the quantity entrants may sell. Unlike a divestiture, the remedy only increases competition if third parties, the potential entrants, find it profitable to enter on those terms, and the authority must rule before those decisions are made, approving the merger whenever entry on those terms would offset the harm to consumers. The effectiveness of the remedy thus depends on decisions the authority does not control, made on terms set by the firms that would lose from the new competition. In light of these features, we ask whether such remedies achieve their stated goal and identify the conditions under which they cannot.

[1] A structural remedy is executed once and is largely irreversible, whereas a behavioral one must be monitored for its whole term (Organisation for Economic Co-operation and Development, 2012; International Competition Network, 2016). Structural remedies account for 85% of FTC merger remedy cases between 2006 and 2012 (Federal Trade Commission, 2017) and 83% of European Commission conditional clearances from 2004 to 2018 (Maier-Rigaud and Loertscher, 2020). See also U.S. Department of Justice, Antitrust Division (2020) and Competition and Markets Authority et al. (2021).

[2] Labeled *access* or *quasi-structural* remedies in the literature. Common examples include obligations to license key technologies, disclose information, relax exclusivity clauses, or grant rivals access to essential inputs, platforms, or production capacity (Motta et al., 2002; Organisation for Economic Co-operation and Development, 2012; Maier-Rigaud and Loertscher, 2020).

[3] This is the case in the European Union (European Commission, 2008) and in many other jurisdictions, including Chile. Exceptions include the United Kingdom, where the authority can impose remedies after an in-depth investigation, and the United States, where remedies are negotiated with the agencies.

We study this question in the context of the 2021 merger between Fiat Chrysler Automobiles (FCA) and Peugeot S.A. (PSA) in Chile. The Fiscalía Nacional Económica (FNE), Chile's competition authority, found that the merger would raise the prices of small commercial vans, a segment in which the two firms jointly accounted for about 60% of sales, and approved it subject to remedies (Fiscalía Nacional Económica, 2020). The central remedy was entry-facilitating and built on an existing relationship. In Europe, PSA already manufactured the Toyota Proace City, a vehicle virtually identical to the merging parties' own models, under a cooperation agreement that did not cover Chile, so Toyota had no product with which to enter the Chilean segment. The remedy extended the agreement to Chile, subject to an annual cap on Toyota's sales proposed by the merging parties themselves, of at most 1,500 units.[4] However, to date, Toyota has not entered the market.

To understand why, we first develop a model of merger review with an entry-facilitating remedy. We build on the framework of Nocke and Schutz (2018, 2025), recently used by Caradonna et al. (2025) to study how potential entry offsets the harm caused by mergers, and add the institutional features described above. The authority reviews the merger in two steps. It first assesses whether the merger harms consumers and, if it does, whether the proposed remedy would undo that harm if the entrant took it up. The merging parties propose a cap on the entrant's sales, and the entrant decides whether to enter only after the merger has been approved. We classify entrants by the effect of their unrestricted entry on consumers. An entrant is *strong* if, entering without restrictions, it would leave consumers at least as well off as without the merger, and *weak* otherwise. With a weak entrant, no cap can satisfy the authority, and the merger does not take place.

Our main result concerns strong entrants, for which the remedy can always be approved but does not always work. Whether it works depends on the merging parties, who choose the cap, and on the entrant, who decides whether to enter. Two questions determine the outcome. First, can the parties propose a cap that the authority accepts but the entrant declines? If so, entry is avoidable. Second, does a workable cap exist, one that satisfies the authority, keeps the merger profitable for the parties in the affected market[5] and makes entry profitable for the entrant?

When entry is not avoidable, the remedy works. (i) If a workable cap exists, the parties propose the smallest cap the authority accepts, and the entrant enters. (ii) If none exists,

[4]The FNE's public report gives the proposed cap only as a range of 600 to 1,500 units; we evaluate the remedy at its upper end (Section 2).

[5]Throughout, the merger's profitability refers to the parties' profits in the market where the remedy applies, not to the transaction as a whole.

the parties abandon the merger. Either way, consumers end up exactly as well off as without the merger, which is all the authority requires. When entry is avoidable, the remedy fails. The parties propose a cap the entrant declines, (iii) even though a workable cap exists (*strategic deterrence*), or (iv) because none does (*structural failure*). The merger goes through, entry never occurs, and consumers end up worse off than without the merger. The failure would disappear if approval were conditioned on realized entry rather than granted on the promise of it. A cap the entrant declines would then secure nothing, so the parties would have to propose one that draws the entrant in or abandon the merger, and consumers would never end up worse off than without it.

We take the model to the data to measure the competitive effects of the merger and to understand why Toyota did not enter. Was Toyota a strong entrant and, if so, was its absence due to structural failure or to strategic deterrence? A difference-in-differences comparison of the parties' prices with those of non-merging vans finds no statistically significant change after the merger. That estimate cannot answer these questions. It nets the parties' cost savings against any increase in market power, and it says nothing about an entrant that never appeared. We therefore estimate a random-coefficients logit model of demand (Berry et al., 1995; Nevo, 2000)[6] together with a model of supply that does not impose how the merged firm sets its prices (Ciliberto and Williams, 2014; Miller and Weinberg, 2017; Michel et al., 2024). Rather than assuming that FCA and PSA either kept competing or fully coordinated their prices from the first month after the merger, we estimate the degree to which the merged firm took into account the profits of its partner's models when pricing its own, which accommodates a gradual integration of the two firms.

We find that the merger harmed consumers. The merged firm partially internalized the profits of its partner's models, which raised the prices of its vans by 4.0% to 5.5% relative to a scenario without the merger, and cost savings offset only about a third of that increase. Toyota, however, was a strong entrant. Had it entered without restrictions, consumers would have been 10.63% better off than without the merger, and even entry capped at 1,500 units, the scenario the authority evaluated, would have left them 6.75% better off. The authority's assessment was therefore correct on its own terms. But Toyota did not enter, and consumers ended up 5.51% worse off than without the merger, while the merging parties' profits in the segment rose by 15.91%.

Toyota's absence reflects a structural failure rather than strategic deterrence. Because the

[6]The model builds on work on differentiated vehicle markets (Verboven, 1996; Goldberg and Verboven, 2001; Grigolon et al., 2018; Wollmann, 2018; Grieco et al., 2024).

cap limited Toyota's sales, its decision to stay out reveals that its cost of entering exceeded what it could earn at the cap. At any smaller cap, which is what the parties would have needed to keep the merger profitable in Chile, Toyota would have earned even less. No cap could therefore satisfy the authority, keep the merger profitable in Chile and attract Toyota at the same time, so the remedy did not fail because of a miscalibrated cap. Conditioning approval on entry, or letting the authority set the cap, would have left consumers at least at their no-merger level rather than 5.51% below it, since a cap Toyota declined would then have secured nothing. At our estimates neither would have brought Toyota in, and the parties would not have merged.

The closest theoretical benchmark is Caradonna et al. (2025), who treat the merger and the entry it may induce as a single equilibrium rather than assessing entry separately.[7] Their main result concerns efficiencies, which reduce a merger's harm but also reduce the entrant's profit. Efficiencies enter our model but are not its subject. We keep their setting and add an authority and a remedy. In their model the only check on a merger is its profitability; in ours an authority reviews it in two stages, clearing it if it does not harm consumers and otherwise accepting or rejecting the remedy the parties propose. The source of entry also differs, since there the merger itself creates the opportunity for entry, while here entry is possible only through the remedy, and the parties choose its scale. As in their model, the parties anticipate the entrant's decision; the authority does not, since it asks whether entry at the proposed cap would clear the standard, not whether entry will follow. The cap lets the parties exploit that gap, which is why the remedy fails when entry is avoidable.

The paper contributes to the theory of merger remedies. That literature separates remedies the authority sets, which may go beyond restoring pre-merger competition (Vasconcelos, 2010), from remedies negotiated under private information about efficiencies (Cosnita and Tropeano, 2009; Cosnita-Langlais and Tropeano, 2012) or about the competitive harm (Lyons and Medvedev, 2007), and from remedies the parties propose. When the parties propose the remedy under a consumer-surplus standard, they concede only what approval requires (Dertwinkel-Kalt and Wey, 2016; Nocke and Rhodes, 2026). In our setting, this is what happens when entry cannot be avoided, as the parties propose the smallest cap the authority accepts. The failure that arises when entry is avoidable has no counterpart there. In both papers the firm that receives the assets makes no decision after approval, even when the divestiture is meant to induce entry, so an approved remedy always materializes. Our entrant decides after approval, with a fixed cost and the option to stay out, so approval credits entry it

[7] Werden and Froeb (1998) explore the question numerically in Cournot and logit Bertrand examples; Caradonna et al. (2025) give the general analysis.

cannot compel. Related work shows that a remedy relying on a third party may not deliver the competition the authority expects (Farrell, 2003; Cabral, 2003). When the authority screens for harm imperfectly, the parties offer the cheapest remedy that gets the merger approved, and that remedy does not necessarily restore consumer surplus (Jara and Rojas, 2026). The closest institutional analogue to conditioning approval on entry is the fix-it-first remedy, where approval waits for the buyer's acceptance (Dertwinkel-Kalt and Wey, 2016).

Empirically, the evaluation of merger effects runs from retrospectives documenting heterogeneous post-merger price effects (Ashenfelter and Hosken, 2010; Ashenfelter et al., 2013) to structural simulations, both in specific industries (Nevo, 2000; Björnerstedt and Verboven, 2016; Ohashi and Toyama, 2017; Montag, 2023) and across many (Weinberg and Hosken, 2013; Bhattacharya et al., 2026). The evaluation of remedies has concentrated on divestitures.[8] Three structural evaluations are closest to ours. Friberg and Romahn (2015) simulate the Carlsberg–Pripps merger in Swedish beer and find that the divestitures dampened the price increase. Delaprez (2026) estimates conduct around a divestiture in US beer and finds that price coordination with the buyer offsets most of its gains. Jara and Rojas (2026) evaluate every feasible divestiture in a Swedish spirits merger and find that the approved one left consumer surplus below the no-merger level. They also show that the merged firm would recover nearly all the profit the remedy takes if the buyer failed and it reacquired the brands. In all three, the buyer acquires the assets, and when the remedy later fails, it does so because of an external shock after the sale, not because the buyer chose to stay out. To our knowledge there is no structural welfare evaluation of an entry-facilitating remedy, where take-up is the designated entrant's choice. We provide one, using the entrant's observed non-entry to bound its entry cost in place of a point estimate.

Methodologically, we estimate post-merger conduct rather than impose it. Flexible conduct frameworks have been used to study coordinated effects and changes in competitive conduct (Ciliberto and Williams, 2014; Miller and Weinberg, 2017; Michel et al., 2024; Delaprez, 2026; Eisenberg et al., 2026). Michel (2016) allows profit internalization to adjust after an ownership change, and Bhattacharya et al. (2026) document at scale that merger simulations diverge from realized outcomes, and identify the assumption that conduct adjusts instantaneously as one source of the gap. We address this limitation by estimating the degree of profit internalization between FCA and PSA in the immediate post-merger period, rather than imposing full internalization from the first month, an assumption that is too restrictive when integration is expected to be gradual.

[8] Reduced-form retrospectives of divestitures include Lagos (2018) for retail gasoline and Osinski and Sandford (2020) for casinos; Federal Trade Commission (2017) reviews the agency's own orders.

2 Institutional Background

The merger was notified to the FNE in March 2020.[9] The FNE's investigation focused on the small commercial van segment, where the merging parties held a combined market share of approximately 60% (Fiscalía Nacional Económica, 2020).[10] In our 2019 sales data PSA held a market share of 42.6% and FCA 17.9%, giving a pre-merger Herfindahl-Hirschman Index of 2,552 and a post-merger increase of 1,527 points (Table 11).

The FNE concluded that the two firms were close competitors, but identified a pronounced asymmetry in substitution patterns. Diversion from FCA models toward PSA models was substantially higher than the reverse, generating strong incentives for the merged entity to raise prices on the FCA vans, sold in Chile under the RAM brand (the V1000 and V700 City), post-merger.[11] Efficiencies claimed by the merging parties, namely economies of scale and procurement synergies, were assessed as unverifiable.

Before turning to the package it approved, the FNE ruled out structural remedies. Neither party manufactured in Chile, so there was no domestic asset to transfer, and the single plant that built the vans served several national markets across segments, which the FNE judged put its divestiture out of proportion to a problem confined to Chilean small commercial vans (Fiscalía Nacional Económica, 2020).

Instead, the package approved in December 2020 centered on an entry-facilitating remedy.[12] The remedy drew on an existing agreement. Under a cooperation agreement between PSA and Toyota Motor Europe (TME), PSA built the Toyota Proace City, a Toyota-badged derivative of its own Citroën Berlingo, Peugeot Partner and Opel Combo, at the same plant

[9] Case F233-2020. FCA and PSA announced their intention to merge in October 2019, creating Stellantis N.V., at completion the fourth-largest automobile manufacturer in the world by sales volume (Stellantis, 2021). The transaction required approval in 22 jurisdictions, six of which imposed remedies. Closing was conditional on competition approvals in a set of jurisdictions that included Chile. Chile's was among the last outstanding, and either party could terminate the agreement if the required approvals had not been obtained by June 30, 2021. The prospectus states that closing without them would in most cases violate applicable law (Fiat Chrysler Automobiles N.V., 2020). Both the European Union and Chile centered theirs on Toyota's role in the small commercial van segment (Global Competition Review, 2021).

[10] Sectoral regulation in Chile classifies commercial vehicles into minibuses, minitrucks, and vans. The FNE's investigation covered the small commercial vans subsegment, which corresponds to a narrower category than the European definition of small light commercial vehicles.

[11] The FNE's consumer survey estimated diversion from FCA to PSA in the range of 30–40%, against diversion from PSA to FCA of at most 10% (Fiscalía Nacional Económica, 2020). We use these figures in Section 6 as an external validation target for the structural demand estimates.

[12] A secondary remedy required FCA and PSA to maintain separate import and wholesale distribution activities for their small commercial van lines in Chile for three years, primarily affecting customer service variables rather than retail prices (Fiscalía Nacional Económica, 2020).

as those models, for sale in Europe.[13] The parties offered to have PSA make available to Toyota an annual capacity of these vans for export to Chile, should Toyota request it, together with administrative certification support, access to its dealer and distribution network, and discounts on the transfer prices of vehicles, parts and components. The FNE judged the Proace City the closest product to the parties' own, since it shares their platform (Fiscalía Nacional Económica, 2020, para. 389), and the merging parties were its sole upstream source.

Since the PSA–TME agreement had been confined to Europe, Toyota had no contractual access to the Proace City for Chile before the remedy, and therefore no product with which to enter the segment.[14] The remedy therefore set the scale at which entry could occur, from a baseline of zero, rather than restricting supply that would otherwise have been available. The FNE's public report refers to the proposed capacity only as a range of 600 to 1,500 units per year (Fiscalía Nacional Económica, 2020, para. 339), and the level in the final commitments is redacted. We evaluate the remedy at the upper end of that range, the level most favorable to entry.

The FNE evaluated the remedy in two steps. It established that the remedy made Toyota's entry possible, since Toyota had no access to the Proace City for Chile without it, and it assessed whether Toyota's entry at the proposed capacity would offset the merger's effect on the segment (Fiscalía Nacional Económica, 2020, paras. 389–396). It did not require a commitment to enter. The option was at Toyota's sole discretion, and the FNE noted that PSA had no certainty Toyota would take up the vehicles (Fiscalía Nacional Económica, 2020, para. 413). Consulted during the investigation, Toyota stated that it would not rule out evaluating entry if demand warranted it and the technical conditions were in place (Fiscalía Nacional Económica, 2020), and the public report contains no assessment of whether its operating profit at that capacity would cover its cost of entering.[15] Although absent from the small commercial van segment, Toyota was already established in adjacent segments in Chile, notably medium commercial vans through the Hiace. The remedy runs for the initial term of the PSA–TME cooperation agreement, whose expiry date is redacted, and any extension of that agreement must be reported to the FNE (Fiscalía Nacional Económica, 2020, para. 384).[16] Toyota had not entered by the end of our sample, and has not entered

[13] PSA and TME had initiated collaboration in medium commercial vehicles in 2012 and extended it to small commercial vehicles in 2019.

[14] The FNE report states that the geographic restrictions of the PSA–TME agreement left Toyota unable to enter the segment in Chile before the remedy (Fiscalía Nacional Económica, 2020, fn. 279 and para. 390).

[15] That the designated entrant was consulted, and that the remedy was built around it, shows the authority treated Toyota's entry as a concrete possibility.

[16] We presume the agreement remains in force, since the cooperation between Stellantis and Toyota Motor Europe in commercial vans has continued since the merger. Stellantis still produces the Proace City, which was updated in 2024,

since.

Toyota's position in Europe was different. The European Commission examined the same transaction and accepted a remedy of essentially identical form.[17] There, however, Toyota already marketed the Proace City across national markets, so the remedy strengthened a competitor already active in the segment. In Chile, Toyota had no such presence, so the same instrument had to induce entry rather than reinforce it. That difference is decisive. Strengthening an active rival requires no participation decision, while inducing entry requires the entrant to judge that the proposed capacity covers the cost of establishing itself in the segment.

3 Theory

We develop a model of merger review with an entry-facilitating remedy. The framework builds on the aggregative games approach of Nocke and Schutz (2018, 2025) and on the merger-and-entry game of Caradonna et al. (2025) (CMS hereafter), adapted to an institutional setting in which a competition authority applies a two-stage review procedure, entry is possible only through the remedy, and the merging parties propose the quantity the entrant may sell. Derivations and proofs of all results are collected in Appendices A and B.

3.1 Environment and Timing

Players. The strategic players are firms 1 and 2 (the merging parties), the non-merging incumbents, and a potential entrant F (Toyota in our application). Let $\mathcal{F}_0 = \{3, \ldots, F-1\}$ denote the non-merging incumbents. A merger between firms 1 and 2 generates efficiency $E \geq 0$, which raises the merged entity's competitive strength.[18] The game is one of *complete information* among the strategic players. There is also a competition authority, which is not a strategic player. It is a rule-applying institution that checks statutory welfare standards at two procedurally distinct stages of the merger review. First, when the firms decide to merge, it makes the usual welfare assessment of the merger proposal. Second, when the parties propose a remedy, it assesses whether the merger, subject to that remedy, clears the same standard. At this second stage it asks whether the remedy makes entry possible and whether

and in 2022 the two companies extended their cooperation to large vans (Stellantis, 2022).

[17] Case M.9730. The Commission required PSA to increase the production capacity available to TME and to reduce transfer prices for vehicles and spare parts (European Commission, 2021). A second remedy addressed access to FCA and PSA repair networks.

[18] Appendix A defines it formally in terms of firm types.

entry at the proposed cap would clear the standard. It requires neither a commitment to enter nor a check that entry would be profitable.

The remedy. The remedy allows the entrant to sell in the market and sets the quantity $\bar{Q}$ it may sell, from a baseline of zero under no remedy. Absent the remedy the entrant cannot enter at all, whatever its fixed cost, as Section 2 documents. We write F_T for the entrant's entry cost net of the remedy's cost-side provisions.[19]

Timing. The sequence of events is as follows.

1. **Merger decision.** Firms 1 and 2 decide whether to merge, anticipating the full continuation of the game. They merge if and only if the profit they realize covers their pre-merger profits:

$$\Pi^{M}_{\text{real}}(\bar{Q}^*, E) \;\geq\; \Pi^{1}_{\text{pre}} + \Pi^{2}_{\text{pre}}, \tag{MP}$$

   where $\bar{Q}^*$ is the equilibrium capacity assigned to the entrant, characterized in Section 3.3, and $\Pi^{M}_{\text{real}}(\bar{Q}^*, E)$ is the merged entity's profit with the entrant present at that capacity if entry occurs, and its profit without the entrant if it does not. If they do not merge, the game ends.[20]

2. **Initial merger screen.** The authority compares consumer surplus under the merger without entry ($\text{CS}_{\text{m,ne}}$) with the no-merger benchmark (CS_0). If the merger passes, it is approved without a remedy and the game ends. If it fails, the merger can proceed only with a remedy.

3. **Cap proposal.** The merging parties propose a capacity $\bar{Q} \geq 0$.

4. **Conditional merger remedy evaluation.** The authority evaluates welfare conditionally, assuming the entrant enters at $\bar{Q}$, and approves if consumer surplus under

[19] The remedy also included homologation support, dealer-network access and preferential transfer prices. The first two lower the fixed cost of entering and are absorbed in F_T. Transfer prices are a per-unit term; the counterfactual assigns Toyota the platform models' production cost (Appendix C), so any upstream margin the parties would earn on Toyota's units is set aside there.

[20] Two ties are resolved as follows. At an approved capacity that leaves profits unchanged, so that (MP) holds with equality, the parties merge. When the parties anticipate that the merger would be blocked, merging and not merging both leave them with their pre-merger profits, and they do not merge.

capped entry clears the same benchmark:

$$\mathrm{CS}(\bar{Q}, E) \;\geq\; \mathrm{CS}_0\,. \tag{A}$$

A proposal that fails (A) is rejected and the merger blocked.

5. **Entry decision.** The entrant enters if and only if its operating profit under the cap covers its entry cost:

$$\pi^F(\bar{Q}, E) \;\geq\; F_T, \tag{F}$$

where $\pi^F(\bar{Q}, E)$ is its profit when it may sell at most $\bar{Q}$. A cap above the entrant's unconstrained quantity does not bind, so π^F equals its unconstrained profit.

6. **Pricing equilibrium.** Active firms compete in prices. The entrant, if it entered, sells at most $\bar{Q}$.[21]

Strategies. A strategy for the merging parties specifies a merger decision, a capacity $\bar{Q} \geq 0$ at each history at which they propose, and a price for each of their products at each history at which they price. A strategy for the entrant specifies, for each approved capacity, whether to enter, and a price at each history at which it has entered. A strategy for each non-merging incumbent specifies a price at each history. The solution concept is subgame perfect equilibrium (SPE).

3.2 Market Payoffs

Four configurations of the market are relevant. They are the no-merger benchmark, the merger without entry, and the merger with entry either uncapped or subject to the cap $\bar{Q}$. Table 1 collects the payoff of each party in each case. We write Q_{unc} for the quantity the entrant would sell if the cap were absent. A cap at or above Q_{unc} does not restrict the entrant, so uncapped entry is simply the special case of capped entry with $\bar{Q} = Q_{\mathrm{unc}}$.

The three constraints introduced above, approval (A), merger viability (MP) and entry (F), are all comparisons within this table. Approval compares the first column's last row with its first. Merger viability does the same in the second column. Entry compares the third

[21] Because products are differentiated, an entrant facing a binding cap prices up until demand equals the cap rather than posting a lower price and turning buyers away, which would forfeit margin on the same quantity. No excess demand arises, so the equilibrium does not depend on a rationing rule.

Table 1: Payoffs by configuration

	Consumers	Merging parties	Entrant
No merger, no entry	CS_0	$\Pi^1_{\mathrm{pre}} + \Pi^2_{\mathrm{pre}}$	—
Merger, no entry	$\mathrm{CS}_{\mathrm{m,ne}}$	$\Pi^M_{\mathrm{m,ne}}$	0
Merger, uncapped entry	$\mathrm{CS}_{\mathrm{unc}}$	$\Pi^M(Q_{\mathrm{unc}}, E)$	π^F_{unc}
Merger, entry at cap $\bar{Q}$	$\mathrm{CS}(\bar{Q}, E)$	$\Pi^M(\bar{Q}, E)$	$\pi^F(\bar{Q}, E)$

column's last row with F_T. The table can also be used to measure what a failed remedy costs consumers. When the authority approves a cap at which entry does not occur, consumer harm is the wedge $\mathrm{CS}_0 - \mathrm{CS}_{\mathrm{m,ne}}$ between the first two rows of the first column.

Appendix A derives these payoffs from multinomial logit (MNL) demand and firm types. Two monotonicity properties, established there from the comparative statics of Nocke and Schutz (2018) (Proposition 6), generate the thresholds of the next section. On the binding region $\bar{Q} < Q_{\mathrm{unc}}$, consumer surplus is strictly increasing in the cap and the merging parties' profit strictly decreasing in it (labeled (M1) and (M3) in Appendix A.3). On the slack region $\bar{Q} \geq Q_{\mathrm{unc}}$ the ceiling is inert and both are constant.[22]

We study the parameter region in which the remedy is the authority's binding instrument, that is, in which the merger can be approved only with a remedy and the choice of cap determines the outcome.

We write $\Pi_0 \equiv \Pi^1_{\mathrm{pre}} + \Pi^2_{\mathrm{pre}}$ for the merging parties' joint pre-merger profit.

Assumption 3.1 (Remedy Necessary)**.** Without a remedy the merger harms consumers: $\mathrm{CS}_{\mathrm{m,ne}} < \mathrm{CS}_0$.

Assumption 3.2 (Cap Necessary)**.** Under uncapped entry the merger is not profitable: $\Pi^M(Q_{\mathrm{unc}}, E) < \Pi_0$.

Absent entry the merger is the static merger of Nocke and Schutz (2025), and by their Proposition 2 it raises the parties' joint profit, $\Pi^M_{\mathrm{m,ne}} > \Pi_0$, for every $E \geq 0$. The two

[22] These properties are established for MNL demand. Our estimated demand system is random-coefficients logit, so the threshold caps and the welfare frontier of Section 7 are computed numerically rather than inheriting the closed-form guarantees. Monotonicity is verified numerically for CS and Π^M at the estimated conduct, which is what Q_A and Q_{MP} require, and Table 13 reads the same thresholds at Nash conduct and at full internalization. It does not extend to the entrant's own profit, which peaks at an interior cap (Proposition A.2).

assumptions therefore place an upper bound on merger efficiency, and both hold at our estimates, as Section 7 reports. Lemma A.1 establishes the bounds $\tilde{E}$ and $\bar{E}$ and their uniqueness, and defines the admissible region $\mathcal{A} \equiv \{E \geq 0 : E < \min(\tilde{E}, \bar{E})\}$ they jointly determine.

3.3 Equilibrium Caps and Entry

A single instrument has to do three jobs at once, one for each party. The cap must be large enough for the authority to approve (A) and small enough for the merger to remain viable (MP), so these two pull in opposite directions. The entrant's constraint (F) pulls both ways. Its operating profit peaks at an interior cap, so entry needs a cap that is not too small and, when F_T exceeds π^F_{unc}, not too large either. Whether any cap satisfies all three at once is settled by E and F_T before any proposal is made.

Each constraint holds on a set of caps.

- *Approvable*, $[Q_A, \infty)$, because surplus rises with the cap, where Q_A solves $\text{CS}(Q_A, E) = \text{CS}_0$. Q_A exists precisely when unconstrained entry would at least offset the merger, $\text{CS}_{\text{unc}} \geq \text{CS}_0$.

- *Merger-viable*, $[0, Q_{MP}]$, because merged profit falls with the cap, where Q_{MP} solves $\Pi^M(Q_{MP}, E) = \Pi_0$. Q_{MP} exists precisely when merged profit is above Π_0 at a zero cap and below it under uncapped entry, as the profitability of the merger without entry and Assumption 3.2 ensure.

- *Entry-inducing*, $[Q_F, Q_F^+]$, because the entrant's profit is single-peaked in the cap, where $Q_F \leq Q_F^+$ solve $\pi^F(Q, E) = F_T$, with $Q_F^+ = \infty$ when $F_T \leq \pi^F_{\text{unc}}$. Q_F exists precisely when $F_T \leq \pi^F_{\max} \equiv \max_{\bar{Q}} \pi^F(\bar{Q}, E)$; otherwise no cap induces entry and the set is empty.

The condition for Q_A to exist classifies the entrant. It is strong when unconstrained entry would leave consumers at least as well off as the no-merger benchmark, $\text{CS}_{\text{unc}} \geq \text{CS}_0$, and weak when $\text{CS}_{\text{unc}} < \text{CS}_0$.[23] Strength depends only on the entrant's competitive effect conditional on entry, not on its fixed cost, so a strong entrant may still stay out. Against a weak entrant no cap is approvable and the merger does not occur (Proposition 3.1), so the discussion that follows concerns the strong entrant.

[23] At the boundary $\text{CS}_{\text{unc}} = \text{CS}_0$, (A) holds with equality at every cap at or above Q_{unc}, so an approvable set exists, with $Q_A = Q_{\text{unc}}$, and we classify the entrant as strong.

A cap is both approvable and merger-viable at entry if and only if it lies in the corridor

$$\mathcal{C} \equiv [\, Q_A,\, Q_{MP} \,],$$

nonempty if and only if $Q_A \leq Q_{MP}$. Merger viability matters only if entry occurs; when entry fails, the merged entity earns its no-entry profit, which clears (MP) at every capacity, since the merger is profitable without entry. Adding the entrant's constraint gives the *success corridor*

$$\mathcal{S} \equiv \mathcal{C} \cap [\, Q_F,\, Q_F^+ \,],$$

the caps that are approvable, merger-viable and entry-inducing simultaneously. We call such a cap *workable*. When entry is unavoidable the entrant's constraint is already implied by the authority's, so $\mathcal{S} = \mathcal{C}$ there.[24]

$\mathcal{S}$ is empty when either set is empty, and otherwise when the two are disjoint. Disjointness takes two forms. Every merger-viable cap may be too tight to induce entry ($Q_F > Q_{MP}$), or every approvable cap may be too generous ($Q_F^+ < Q_A$), the second possible because π^F is single-peaked rather than monotone.

The proposed cap follows from the parties' payoff, and two facts rank the candidates. Because Π^M is strictly decreasing in $\bar{Q}$, among approvable capacities that induce entry the smallest is best for the parties. But a capacity that deters entry pays the no-entry profit, which exceeds $\Pi^M(\bar{Q}, E)$ at every capacity that induces entry. The parties therefore propose an approvable capacity that deters entry whenever one exists. Only when no such capacity exists do they solve

$$\bar{Q}^* \in \arg\max_{\bar{Q} \geq 0} \; \Pi^M(\bar{Q}, E) \quad \text{s.t.} \quad \text{(A)}, \tag{1}$$

whose solution is Q_A.

Entry can be deterred in two ways: with a cap too tight to cover the fixed cost, or, because π^F peaks at an interior cap, with one too generous. It is unavoidable exactly when neither kind is approvable. $Q_F \leq Q_A$ rules out the first and $F_T \leq \pi^F_{\text{unc}}$ the second.

Two questions therefore settle the outcome. Can the parties avoid entry, and does a workable

[24] Both endpoints move with efficiency. A more efficient merger clears the standard at a smaller cap and stays viable at a larger one, so $\mathcal{C}$ widens in E (Lemma B.1), while the entry-inducing set narrows, since efficiency lowers the entrant's profit (Proposition A.2). The effect on $\mathcal{S}$ is therefore ambiguous.

cap, one in $\mathcal{S}$, exist? Taken together, the answers give the four cases below.

Proposition 3.1 (Equilibrium Cap and Entry)**.** In every SPE, prices at each pricing history are the unique Bertrand equilibrium of Appendix A, and the entrant enters if and only if $\bar{Q} \in [Q_F, Q_F^+]$. Let $\mathcal{D} \equiv [Q_A, \infty) \setminus [Q_F, Q_F^+]$ be the approvable capacities that deter entry. The equilibrium outcome is unique. If the entrant is weak, no capacity is approvable, the parties do not merge, and $\mathrm{CS} = \mathrm{CS}_0$. If it is strong, one of four cases obtains:

(i) **Minimal entry.** Entry unavoidable and $\mathcal{S} \neq \emptyset$. The parties merge and propose $\bar{Q}^* = Q_A$. The entrant enters, and $\mathrm{CS} = \mathrm{CS}_0$.

(ii) **Merger abandoned.** Entry unavoidable and $\mathcal{S} = \emptyset$. The parties do not merge, and propose some $\bar{Q}^* < Q_A$. The entrant cannot enter, and $\mathrm{CS} = \mathrm{CS}_0$.

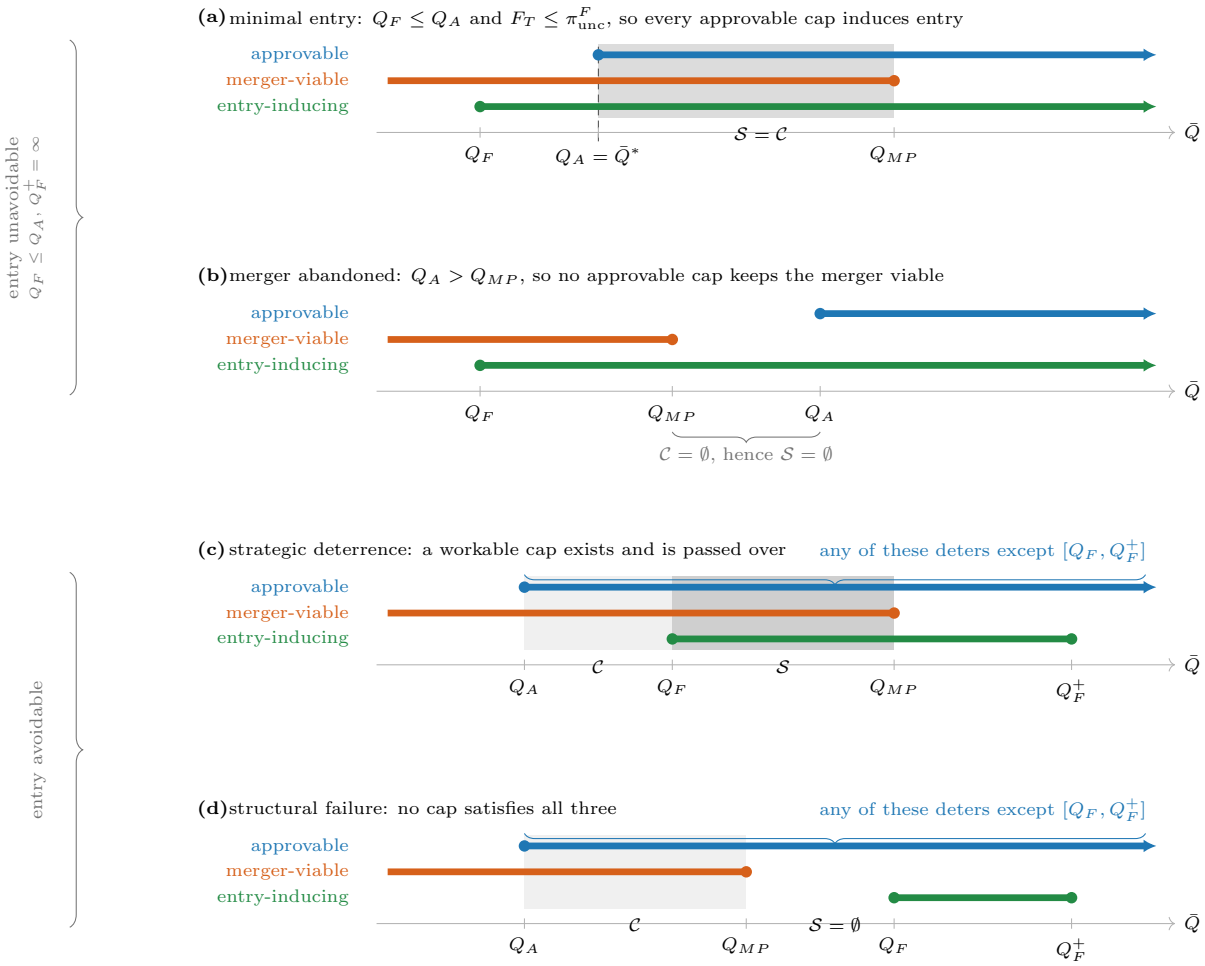
(iii) **Strategic deterrence.** Entry avoidable and $\mathcal{S} \neq \emptyset$. The parties merge and propose some $\bar{Q}^* \in \mathcal{D}$. The entrant stays out, and $\mathrm{CS} = \mathrm{CS}_{\mathrm{m,ne}} < \mathrm{CS}_0$.

(iv) **Structural failure.** Entry avoidable and $\mathcal{S} = \emptyset$. The parties merge and propose some $\bar{Q}^* \in \mathcal{D}$. The entrant stays out, and $\mathrm{CS} = \mathrm{CS}_{\mathrm{m,ne}} < \mathrm{CS}_0$.

The equilibrium proposal is unique in case (i). Otherwise the set of equilibrium proposals is $[0, Q_A)$ in case (ii), $\mathcal{D}$ in cases (iii) and (iv), and $[0, \infty)$ against a weak entrant. Within each set all proposals yield the same outcome.

Consumer surplus depends only on whether entry is avoidable. It equals CS_0 in cases (i) and (ii) and $\mathrm{CS}_{\mathrm{m,ne}}$ in (iii) and (iv). Whether a workable cap exists plays a different role in each case. When entry is unavoidable, it determines whether the merger takes place, since an empty success corridor means that every approvable capacity leaves the parties with less than Π_0, so they do not merge. When entry is avoidable, the parties merge and deter entry regardless, and the corridor determines only whether a capacity existed that would have induced entry while keeping the merger viable.

Figure 1: The three constraint sets and the cases of Proposition 3.1



Notes: each panel places the three constraint sets on the capacity axis. Light shading is the corridor $\mathcal{C} = [Q_A, Q_{MP}]$, darker shading the success corridor $\mathcal{S}$ of workable caps. Arrowheads mark sets unbounded above and filled dots closed endpoints. Panels (a)–(d) are cases (i)–(iv) of Proposition 3.1. Panel (d) draws a finite Q_F; when $F_T > \pi^F_{\max}$ the entry-inducing set is empty and the case is unchanged.

3.4 What the Remedy Achieves

Merger control asks a remedy to undo the merger's harm, leaving consumers at least as well off as without the merger. Under the procedure the authority applied, the remedy meets this standard only when entry cannot be avoided. Realized surplus equals CS_0 against a weak entrant and in cases (i) and (ii), and falls to $\text{CS}_{\text{m,ne}} < \text{CS}_0$ whenever entry is avoidable, in cases (iii) and (iv) (Proposition 3.1). Even when it meets the standard, the remedy never delivers more than CS_0, so consumers are weakly worse off facing a strong entrant rather than a weak one, which is the sense in which the remedy does not reward strength. This

subsection asks which feature of the procedure is responsible, by changing one element of the game at a time and recording what consumers receive.

Against a strong entrant, the type a competition authority seeks, the remedy thus at best meets the standard and at worst fails. Neither outcome reflects a lenient authority. Both arise because the terms of entry are set by the party that gains from limiting it, and that party holds two levers. Because it chooses the cap, it concedes no more than approval requires. When entry cannot be avoided and a workable cap exists (case (i)), it proposes the smallest cap the authority accepts, and the entrant's procompetitive value accrues to the parties rather than to consumers. Because approval does not depend on whether entry occurs, it can also propose a cap the entrant declines, and it does so whenever entry is avoidable (cases (iii) and (iv)), since deterrence pays more than any capacity that induces entry. Among the regimes below, verified entry removes the second lever and an authority-set cap removes both.

Proposition 3.2 (Institutional Regimes)**.** Regime (b) is the baseline game. Each other regime changes one element of it, stated in italics. In every regime prices are as in Proposition 3.1, the entrant enters if and only if (F) holds at the capacity in force, and the parties merge only if (MP) holds there.

(a) **No remedy.** *The remedy stage is removed.* The parties do not merge. The entrant cannot enter, and $\mathrm{CS} = \mathrm{CS}_0$.

(b) **Party-proposed cap** (the baseline). By Proposition 3.1, $\mathrm{CS} = \mathrm{CS}_0$ against a weak entrant and when entry is unavoidable, and $\mathrm{CS} = \mathrm{CS}_{\mathrm{m,ne}} < \mathrm{CS}_0$ when it is avoidable.

(c) **Uncapped access.** *The proposal stage is removed and the capacity is unbounded.* Against a weak entrant, or a strong one with $F_T \leq \pi^F_{\mathrm{unc}}$, the parties do not merge, the entrant cannot enter, and $\mathrm{CS} = \mathrm{CS}_0$. Against a strong entrant with $F_T > \pi^F_{\mathrm{unc}}$ the parties merge, the entrant stays out, and $\mathrm{CS} = \mathrm{CS}_{\mathrm{m,ne}} < \mathrm{CS}_0$.

(d) **Verified entry.** *Approval is conditioned on realized entry.* If $\mathcal{S} \neq \emptyset$ the parties merge and propose $\bar{Q}^* = \min \mathcal{S} = \max(Q_A, Q_F)$. The entrant enters, and $\mathrm{CS} = \mathrm{CS}(\bar{Q}^*, E) \geq \mathrm{CS}_0$, strictly above it if and only if $Q_F > Q_A$. If $\mathcal{S} = \emptyset$ the parties do not merge, the entrant cannot enter, and $\mathrm{CS} = \mathrm{CS}_0$.

(e) **Authority-set cap.** *The authority selects the capacity, maximizing consumer surplus subject to* (A)*,* (MP) *and* (F)*.* If $\mathcal{S} \neq \emptyset$ it selects $\bar{Q}^* = \max \mathcal{S}$, the parties merge, the

entrant enters, and $\text{CS} = \text{CS}(\max \mathcal{S}, E)$. If $\mathcal{S} = \emptyset$ the parties do not merge, the entrant cannot enter, and $\text{CS} = \text{CS}_0$.

Comparing each regime with the baseline shows what each change achieves, and Proposition B.4 ranks all five. Removing the remedy returns consumers to CS_0, so the baseline never does better than no remedy, and does strictly worse whenever entry is avoidable. Uncapped access does no better, since an entrant that would enter at unbounded scale makes the merger unprofitable, while one that would not lets the merger go through without entry, leaving the parties their no-entry profit. Verified entry makes approval conditional on entry actually occurring, so a cap the entrant declines no longer secures the no-entry profit, and consumers are never left below CS_0. They can also end up better off. When a workable cap exists and inducing entry requires more capacity than the authority demands ($Q_F > Q_A$), the parties must propose at least Q_F, and consumers keep the additional surplus that this larger capacity generates. When $Q_F \leq Q_A$ they receive exactly CS_0, and when no workable cap exists the parties do not merge. Only an authority-set cap passes the entrant's full value on to consumers whenever a workable cap exists, because it takes the choice of scale away from the parties.

The two improvements differ in what they require. Verified entry needs approval that can be withheld or reversed when entry fails, and access that is monitored rather than assumed. That is more demanding than a divestiture, whose sale completes once. An authority-set cap needs the same and, in addition, the entrant's fixed cost and the parties' profits at entry, so we report it as a bound rather than a recommendation. Protecting consumers with an entry-facilitating remedy is therefore attainable, but not when approval credits entry before it occurs. The remedy's failure is a property of unverified approval, rather than of capacity remedies as such or of the consumer welfare standard itself. Approval rests on entry credited but never checked, so a merger that harms consumers clears on the strength of an entrant that never appears. Section 7 takes the classification to the data.

4 Data and Reduced-Form Evidence

4.1 Data

All vehicles sold in Chile are imported, and several layers of agents sit between manufacturer and consumer. However, the manufacturer's wholesale price is the key determinant of the suggested retail price and changes in it are largely passed through (Fiscalía Nacional

Económica, 2020), which is what lets us model manufacturer pricing directly. Small commercial vans form a distinct segment. They are productive inputs bought by small firms and contractors for last-mile logistics, so buyers weigh total cost of ownership heavily and are more price-sensitive than passenger-vehicle consumers (Department for Transport, 2024; European Commission, 2021).

Our main data set covers prices, sales and physical characteristics of every small commercial van model sold in Chile between January 2015 and March 2023, following Gómez-Lobo and Price (2025). Daily retail sales come from National Treasury (TGR) records, linked through the Technical Report Code (CIT) to technical specifications registered by the Ministry of Transport. Cost shifters come from Central Bank real exchange rates, FRED steel price indices, and transaction-level customs import records giving CIF and FOB prices, freight and insurance costs, and country of origin. We homologate rebadged models and define firms at the level of economic groups rather than commercial brands, so PSA's Peugeot, Opel and Citroën offerings are a single firm. Sample construction, the time window, model turnover and summary statistics are in Appendix F (Tables 9 and 10).

The dataset contains 83,198 transactions, an average of 840 vans per month, which aggregate to 2,248 year-month-product observations across 14 manufacturers and 52 models. The segment was already concentrated, and the merger raised the HHI to 4,079. PSA and FCA were the largest sellers, followed by Chevrolet and Renault, with every other participant below 5% (Table 11). Their combined monthly share moved little afterwards, from 60% to 58%. Both the number of firms and the number of products offered by non-merging firms decline toward the end of the sample (Table 9). The decline began before the merger was announced and consists of small sellers leaving the segment and of models replaced by their successors. Only one firm enters over the full period, Brilliance, with the X30 in July 2021, at negligible share. Post-merger, then, the merging parties' share was stable and there was essentially no entry.

The segment is divided by place of manufacture. About 30% of the vans sold are made in China, and they form a lower price tier. Their prices barely overlap with those of vans made elsewhere, mostly in Europe, and none of them has a diesel engine (Table 2). Observed characteristics explain only part of this price gap. In a regression of log price on an indicator for manufacture outside China, the characteristics that enter the demand model and year-month fixed effects, vans made outside China carry a premium of 0.37 log points (SE 0.05), compared with 0.60 when the characteristics are excluded. Restricting the sample to gasoline vans, which holds engine type fixed, the premium is 0.37 (SE 0.04), with 5 of these models

made outside China. These results reproduce the FNE's finding that buyers pay significantly more for European vans after controlling for observable characteristics, and its view that origin and engine type are what differentiate light vans in Chile (Fiscalía Nacional Económica, 2020). The premium reflects attributes the data do not record, with origin serving as their proxy. Moreover, because no China-made van is diesel, engine type differentiates vans only among those made outside China, whereas origin differentiates the whole segment.

Table 2: Small commercial vans by place of manufacture

	Made in China	Made outside China
Share of units sold (%)	29.5	70.5
Products	17	35
Price (thousands of CLP)	6,104	9,764
Price, 10th percentile	4,757	8,017
Price, 90th percentile	7,795	12,081
Diesel (0–1)	0.00	0.89
Weight (kg)	1,856	2,092
Wheelbase (mm)	2,730	2,713
Engine size (L)	1.31	1.49
Performance (km/liter)	11.2	15.9

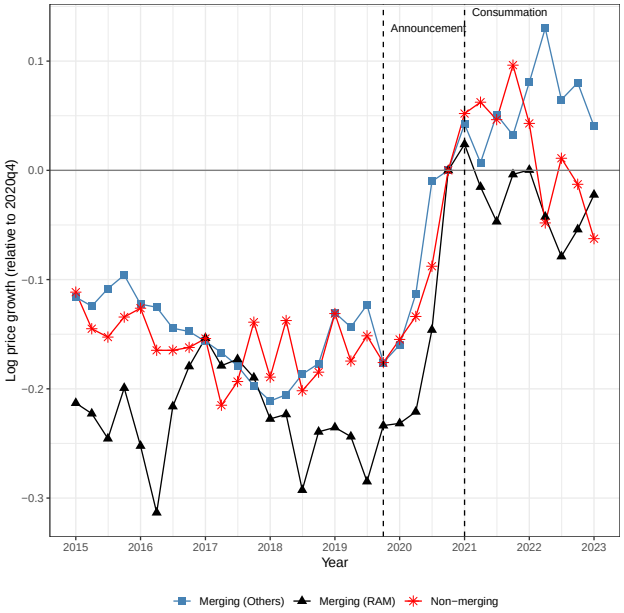
Notes: January 2015 to March 2023, weighted by units sold. Place of manufacture, not the brand's home country, so the Chevrolet N300 and N400 count as made in China. Price percentiles are sales-weighted.

4.2 Reduced-Form Evidence

Figure 2 plots normalized prices for three groups: the RAM models, the remaining merged products (PSA's Partner, Berlingo and Combo, plus the RAM V700 Rapid), and non-merging products.[25] Relative to the last pre-merger quarter, RAM prices are flat, while the other merged products rise faster than non-merging products. That faster rise is a composition effect, driven by PSA's replacement of its previous generation of models, which began before the merger was announced. Once the set of products is held fixed, the other merged products track non-merging prices after consummation, as do the RAM models (Appendix S1).

[25]The RAM V700 Rapid, being gasoline-powered, was a more distant substitute for PSA's diesel offerings and raised no competitive concern in the FNE's assessment.

Figure 2: Average log price over quarterly intervals



Notes: Normalized to 2020q4, the last quarter before consummation.

We estimate a difference-in-differences regression of $\log p_{jt}$ on $\text{Merged}_j \times \text{Post}_t$ with product and month fixed effects, clustering at the product level. Applying this design to a national merger is difficult. An unaffected comparison group is hard to identify and estimates are sensitive to how it is chosen (Friberg and Romahn, 2015; Taylor and Hosken, 2007; Houde, 2012). Following the reduced-form merger literature (e.g., Ashenfelter and Hosken, 2010; Weinberg and Hosken, 2013; Ohashi and Toyama, 2017) we use non-merging small vans as the control, which shares demand and cost characteristics with the FCA-PSA products, and we use two treatment dates, October 2019 (announcement) and January 2021 (consummation), to bound the onset of effects.

The estimates are small, of either sign, and statistically insignificant. That holds in the pooled merged portfolio, in the RAM and non-RAM groups separately, under either treatment timing (Table S1), and with sales weights.[26] With treatment dated at consummation, the pooled estimate is 0.015 (SE 0.027) and the estimate for the RAM group is 0.021 (SE 0.030).

The absence of a significant price effect cannot tell us whether the remedy worked. The design measures the net effect of the merger on the parties' prices, in which the exercise of

[26] Weighting observations by sales, the pooled coefficient at consummation is 0.001 (SE 0.024), the RAM group 0.019 (0.024) and the remaining merged products −0.011 (0.030). The full set of weighted estimates is available on request.

market power is partly offset by the parties' own cost savings. A small net effect is therefore compatible with a substantial increase in market power, which is why we turn to a structural model that separates the two.

5 Structural Model

Post-merger price movements can reflect shifts in demand, in marginal costs, or in conduct, and the reduced-form design cannot separate them. We therefore estimate a structural model of demand and supply for small commercial vans that recovers each channel, and in particular the degree to which the merged entity internalizes the profits of its partner's models.

5.1 Demand

We estimate a random-coefficients logit model of demand in which consumers differ in how they value manufacture outside China. Flexible substitution patterns matter here because own- and cross-price elasticities are the primary determinants of the pricing first-order conditions from which we recover conduct. Consumers in month t choose one of the J_t small commercial vans on offer or an outside option (a second-hand van, a minitruck, or another transport mode). Mean utility for van j is

$$\delta_{jt} = X_{jt}\beta - \alpha p_{jt} + \xi_{jt}, \tag{2}$$

where X_{jt} is a $1 \times K$ vector of product characteristics, β a $K \times 1$ vector of taste parameters, p_{jt} the price of product j, α the price coefficient, and ξ_{jt} an unobserved demand shock. Mean utility for the outside good is normalized to zero. Consumers differ in a single dimension,

$$\mu_{ijt} = \sigma_o \, o_j \, \nu_i \tag{3}$$

where o_j indicates manufacture outside China, predominantly in Europe, and ν_i is an i.i.d. standard normal taste shock. With an i.i.d. Type I Extreme Value shock completing utility, shares take the standard mixed-logit form, integrated over ν_i. Origin is where we let tastes differ because it is what divides the segment into two price tiers that observed characteristics

do not explain (Section 4). Because o_j also enters X_{jt}, β carries the mean taste for origin and σ_o its dispersion across consumers; the price coefficient α is common across consumers. Section 6 assesses this specification against external information from the FNE's consumer survey.

Since p_{jt} may be correlated with the unobserved characteristics ξ_{jt} in (2), as firms set prices knowing product- and market-specific consumer valuations, we instrument with cost shifters: the interaction between international steel prices and each model's weight, entered in squares, and the per-vehicle price of import freight. The maintained exclusion restriction is that, conditional on observed characteristics and year and month-of-year fixed effects that absorb aggregate demand movements, these shift costs affect demand only through prices. Identification comes from differential exposure across models, through weight and origin, to common input-price and freight-cost movements.

We also need instruments for the nonlinear parameter. Following Gandhi and Houde (2019), we construct differentiation instruments (Diff IVs): functions of distances in observable characteristics space between each model and the other models offered by its own firm and by rivals, weighting nearby models more heavily.[27] For σ_o the instruments that matter are those built on origin: the number of rival models made in the same place and rivals' proximity in origin. As the mix of same-origin rivals changes across months, how shares respond reveals how strongly buyers sort by origin, which separates σ_o from the mean taste. Validity requires that, conditional on observed characteristics and fixed effects, the product set and characteristics, and hence the proximity measures, are predetermined with respect to the within-market demand shocks ξ_{jt}.

Estimation follows the standard procedure. For a candidate parameter vector $\theta^D = (\beta, \alpha, \sigma_o)$ we invert observed shares for mean utilities using the Berry et al. (1995) contraction, recover the demand shocks $\xi_{jt}(\theta^D)$ from (2), and choose θ^D to minimize a GMM objective built on the orthogonality conditions $E[\xi_{jt}(\theta^D) \mid Z_{jt}] = 0$. Estimation proceeds in two stages, the second using approximated optimal instruments constructed from the first-stage estimates (Reynaert and Verboven, 2014; Conlon and Gortmaker, 2020), and we report the second-stage estimates. Appendix E states the moment conditions and the algorithm in full. We define the number of potential consumers M_t as recorded sales scaled by $(1 + r)$ and set $r = 0.5$.[28]

[27] The baseline builds them under the ownership in force each year, so FCA and PSA models count as one firm's products from 2021.

[28] Following Ivaldi and Verboven (2005) we also report $r = 1$ and $r = 2$, a potential market 100% and 200% larger than recorded sales. The assumption shifts the outside-good share and hence the level of mean utilities recovered from the

5.2 Supply

On the supply side, we model price competition among differentiated products under Bertrand–Nash behavior, allowing for partial or full internalization of multiproduct pricing externalities between FCA and PSA in the post-merger period. In this setting, each firm f sets the prices of the products $\mathcal{J}_{ft}$ it offers in market t to maximize the weighted profit

$$\pi_{ft} = \sum_{j \in \mathcal{J}_{ft}} (p_{jt} - c_{jt})\, s_{jt}\, M_t \;+\; \kappa_t \sum_{k \in \mathcal{J}_{g(f)t}} (p_{kt} - c_{kt})\, s_{kt}\, M_t, \tag{4}$$

where p_{jt} denotes the price of product j, c_{jt} its marginal cost, s_{jt} its market share, M_t is the market size, and $g(f)$ is f's merger partner in the post-merger period. The second term is therefore present only for FCA and PSA after the merger and is absent otherwise, so every other firm maximizes its own profit alone. Taking first-order conditions with respect to p_{jt} and canceling the common factor M_t yields

$$s_{jt} + \sum_{k=1}^{J_t} \Omega_{jk,t}(p_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial p_{jt}} = 0, \qquad \forall j, \tag{5}$$

where the ownership matrix Ω_t generalizes the usual matrix of zeros and ones. Here $\Omega_{jk,t} = 1$ if products j and k are owned by the same firm in market t and 0 otherwise, except that $\Omega_{jk,t} = \kappa_t$ whenever product j belongs to FCA and product k to PSA, or vice versa. Setting $\kappa_t = 0$ recovers independent pricing between the merging parties and $\kappa_t = 1$ full internalization.

Intuitively, when products are close substitutes, a price increase for an FCA model diverts sales toward PSA models. With higher κ_t, the merged entity places greater weight on these diverted profits when setting prices, bringing observed pricing closer to the joint-profit benchmark.

We estimate κ as a free parameter rather than imposing $\kappa = 1$ post-merger, following Michel (2016). How far a merged firm internalizes the diversion between its lines in the years after a transaction is an empirical question, and imposing full internalization is one reason merger

inversion, leaving the price coefficient, the implied elasticities, and $\hat{\kappa}$ stable across the three (Tables S3 and S5).

simulations overestimate price changes in the large-scale study by Bhattacharya et al. (2026).

Stacking these conditions gives $s_t(p_t) + \Delta_t(p_t)(p_t - c_t) = 0$, where $\Delta_t(p_t) = \Omega_t \circ (\partial s_t(p_t)/\partial p_t')'$ collects demand slopes and pricing internalization, and p_t and c_t are the vectors of prices and marginal costs. To complete the supply side, we specify marginal costs as

$$c_{jt} = w_{jt}\gamma + \omega_{jt}, \tag{6}$$

where w_{jt} includes product characteristics, the price of freight, linear and quadratic trends, and the interaction of freight cost with post-merger merged ownership, which carries the merger-specific efficiency, while ω_{jt} represents unobserved cost shocks. Because marginal costs are not directly observed, they are recovered from the pricing first-order conditions in equation (5). Substituting the implied markups yields the structural cost residual

$$\omega_{jt} = c_{jt} - w_{jt}\gamma = p_{jt} - w_{jt}\gamma + \left[\Delta_t(p_t)^{-1} s_t(p_t)\right]_j. \tag{7}$$

We estimate marginal costs and the conduct parameter taking the demand estimates $\widehat{\theta}^D$ as given. The markup component in equation (7) is endogenous, since unobserved cost shocks are embedded in prices, so we estimate by GMM on the moment condition $E[\omega_{jt}(\gamma, \kappa)\, W_{jt}] = 0$, where W_{jt} collects the cost shifters in w_{jt} and the Diff IVs. The Diff IVs identify κ, following Michel et al. (2024), Eisenberg et al. (2026) and Delaprez (2026). They enter the implied markup through the demand derivatives in the pricing first-order conditions, and κ scales the cross-product terms in those conditions, so variation in competitive proximity moves predicted markups differently under alternative values of κ. Validity requires that, conditional on observed characteristics, fixed effects and the included cost shifters, local crowding in characteristics space affects prices through markups rather than through unobserved cost shocks ω_{jt}. Separating conduct from merger-specific cost changes also requires a restriction on how those changes enter marginal cost. In the baseline they operate only through freight costs, and Section 6 reports the estimates and varies the restriction.

We form the sample analogue of the moment condition and estimate (γ, κ) by two-step GMM. Appendix E gives the moments and objective.

6 Empirical Results

This section reports the estimates, demand first and then the supply side.

6.1 Demand

Table 3 reports three specifications: a standard logit as a benchmark, a nested logit with nests on origin, and the random coefficient on origin we adopt. It gives the price coefficient, the nesting and random-coefficient parameters, and the implied average own-price elasticity, which are the objects the counterfactuals use. Table S2 reports every characteristic coefficient, together with two specifications without origin structure that we discuss below.

In the logit, all coefficients are precisely estimated and have the expected signs (Table S2). Demand slopes down, and consumers prefer heavier vans with larger engines and better performance, diesel over gasoline, and manufacture outside China. The implied average own-price elasticity is -7.139.

Table 3: Demand estimates

	Logit	NL (origin)	RC (origin)
Price	-0.844^{***}	-0.274^{***}	-0.533^{***}
	(0.059)	(0.060)	(0.038)
ρ (nesting)		0.571^{***}	
		(0.082)	
σ_{origin}			10.820^{***}
			(2.917)
Mean Own Price Elast	-7.139	-5.236	-4.416
Market size	1.5× sales	1.5× sales	1.5× sales

Notes: Price is measured in millions of CLP. All estimations include year and month-of-year FE. The log-normal and one-nest specifications are in Table S2.

The origin random coefficient is large and precisely estimated, $\hat{\sigma}_o = 10.820$ with a standard error of 2.917, so consumers differ substantially in how much they value manufacture outside China. The other coefficients keep their logit signs except one: the mean origin coefficient turns negative once the random coefficient absorbs the heterogeneity it had been standing in for. That negative mean is net of the observed characteristics, which favor vans made

outside China, and even then a sizable share of buyers values that origin positively once their individual taste is added to the mean. The average own-price elasticity is -4.416, with a median of -4.351 and a 10th–90th percentile span of -6.271 to -2.492.

The magnitude of $\hat{\sigma}_o$ is best read as market segmentation rather than as a willingness to pay. Tastes for origin are so dispersed that, holding other characteristics fixed, most buyers strongly prefer one origin over the other,[29] and 34% of all buyers prefer manufacture outside China.[30] A random coefficient this dispersed on a binary characteristic behaves like a nest that groups products by origin (Grigolon and Verboven, 2014), and the nested logit on origin captures the same pattern with $\rho = 0.571$.

The FNE's consumer survey offers an external check on these substitution patterns. It reports 30–40% diversion from FCA to PSA and at most 10% in the reverse direction, with a sampling error of roughly 5% (Fiscalía Nacional Económica, 2020). The random coefficient on origin reproduces the pronounced asymmetry the survey documents. Model-implied diversion from FCA to PSA averages 48.7% over 2015–2020, against 28.1% in the reverse direction. Both levels are higher than the survey's, but the ordering, which drives the merged firm's incentive to raise RAM prices, is the same. Since every specification in Table 4 also reproduces this asymmetry, the survey is more informative about how much substitution each sends toward China-made vans. In it, 82% of FCA buyers and more than 90% of PSA buyers said the van they bought was not, or hardly, replaceable with a China-made one, and only 17% and 12% had requested a quote for one (Fiscalía Nacional Económica, 2020, para. 127). Both specifications with origin structure are consistent with these figures. The random coefficient on origin sends only 8.4% of the sales diverted from FCA and 9.1% of those diverted from PSA to China-made vans, and the nested logit on origin 9.5% and 10.3%, whereas specifications without origin structure send roughly a third.

Flexibility along other dimensions does not substitute for origin structure, as two alternative specifications show. The first adds a log-normal random coefficient on price, paired with one on the constant, and the second is a nested logit with all vans in a single nest (Table S2). The first finds no dispersion in price sensitivity to exploit, $\hat{\sigma}_{\text{price}} = 0.23$ with a standard error of 3.4. The second estimates $\hat{\rho} = 0.97$, which makes every van a close substitute for every other, the opposite of a segment divided by origin. Both send 32.9% and 35.0% of diverted

[29] One standard deviation of origin taste is worth 20.3 million CLP against an average price of 8.7 million, and for 70% of consumers the origin term exceeds the utility of an average van's price in absolute value.

[30] The share of buyers who prefer manufacture outside China, 34%, is below the share of vans sold that are made there, seven in ten (Table 2). Buyers choose on total utility. Vans made outside China are heavier and mostly diesel, and for many buyers these characteristics outweigh both their preference for China-made vans and the 1.60 price ratio.

sales to China-made vans under the first and 33.0% and 35.5% under the second, above the share of FCA and PSA buyers who had requested a quote for one.

Table 4: Substitution patterns across demand specifications

Demand specification	Origin structure	FCA → CHINA	PSA → CHINA	FCA → PSA	PSA → FCA
RC on price (log-normal)	no	32.9%	35.0%	41.9%	23.5%
Nested logit (one nest)	no	33.0%	35.5%	42.4%	23.8%
Nested logit (origin)	yes	9.5%	10.3%	46.7%	26.8%
RC on origin (baseline)	yes	8.4%	9.1%	48.7%	28.1%

Notes: model-implied diversion ratios, computed from the estimated demand parameters of each specification over pooled 2015–2020. Origin structure records whether the model lets consumers value European relative to Chinese manufacture differently, either through a random coefficient on origin or through a nest on it. FCA → CHINA is the share of sales diverted from FCA that goes to Chinese-manufactured vans, and likewise for PSA. Section 6 compares these with the FNE's consumer survey.

Between the two specifications with origin structure, we retain the random coefficient as our baseline. It lets the strength of the preference for manufacture outside China vary across consumers rather than summarizing it in a single within-nest correlation, and the choice matters. The nested logit implies a mean own-price elasticity of -5.236 against -4.416 in the baseline (Table 3). It yields a similar merger harm and also finds no workable cap, but it classifies Toyota as a weak entrant, so only the random coefficient rationalizes the three observed actions of the authority, the merging parties and Toyota (Section 7). Engine type, the other dimension the FNE emphasizes, enters every specification as a product characteristic. As for the elasticities themselves, they are consistent with external evidence. The average elasticity of -4.416 falls inside the band of roughly 3 to 7 typically found for passenger vehicles (Berry et al., 1995; Goldberg and Verboven, 2001), and evidence on commercial vehicles and on South American markets points to comparable or greater price sensitivity.[31]

The margins reported to the FNE offer a second check, once the difference in objects is accounted for. Since the model treats the manufacturer as setting the retail price, our markups measure the margin of the whole distribution chain, from the manufacturer's cost to the price paid by the final buyer. The figures reported to the FNE capture only the

[31] Verboven (1996) and Ohashi and Toyama (2017) report estimates in the same band. Wollmann (2018) finds more elastic demand for commercial vehicles. For the region, Brambilla (2005) reports 3.4 and 1.7 for passenger vehicles in Argentina and Brazil, while Chaves (2022) and Cisternas (2018) document elasticities often exceeding 8.

first link of that chain, the manufacturer's margin on its sales to importers, and exclude the margins of the importers and dealers that stand between manufacturer and consumer (Section 2), so they are a lower bound for ours. Model-implied 2019 markups for the PSA models lie between 30% and 44%, for the RAM V700 City between 27% and 32%, and for the RAM V1000 between 22% and 25%, consistent with the pre-merger averages in Table 6. The reported manufacturer margins are 10–20% for PSA, 0–10% for the RAM V700 City and 20–30% for the RAM V1000 (Fiscalía Nacional Económica, 2020), below ours for PSA and the V700 City and overlapping for the V1000.

Appendix S2 reports the baseline under market sizes of two and three times recorded sales. For the origin random coefficient, the differentiation instruments deliver a Montiel Olea–Pflueger effective F (Montiel Olea and Pflueger, 2013) of 19.4 for $\hat{\sigma}_o$, computed holding $\hat{\alpha}$ at its estimate, above the 5% critical value of 15.2 for a worst-case bias of 10%.

6.2 Supply

On the supply side, we estimate a profit-internalization matrix characterized by a single conduct parameter capturing the degree of coordination between FCA and PSA in the post-merger period.[32] We also allow for merger efficiencies related to discounts in the price of freight, which is relevant for the automobile market (Coşar et al., 2018), especially when available products are all imports.

Table 5 reports the GMM estimates of the conduct parameters and merger-specific efficiency effects, with the full cost-shifter set in Table 8. The estimated post-merger conduct parameter is $\hat{\kappa} = 0.394$ (SE 0.193), which rejects both Bertrand–Nash and full internalization at the 5% level. The 95% interval is [0.02, 0.77]. This result indicates partial profit internalization over the two years and a quarter following the transaction. The level coefficient on freight price is positive and precisely estimated, consistent with all vans in the segment being imported and therefore exposed to transportation and logistics costs, the channel through which the baseline allows a merger-specific efficiency to operate. The freight–ownership interaction is negative for the merging parties, implying a mean marginal-cost reduction of 3.4%, in line with Ohashi and Toyama (2017).

Robustness checks for the conduct estimates are reported in Appendix S3. Across them, the point estimate of κ remains between independent pricing and full coordination. Flexible time controls reduce the precision of the estimate, because they absorb some of the time-series

[32] We take January 2021 as the merger date, corresponding to the consummation of the transaction.

variation that helps identify conduct. Separating κ from merger-specific cost changes rests on the restriction that those changes operate through freight, and Table S7 reports what that restriction delivers. Relaxing it, by adding a post-merger level shift to the merged products' marginal cost, raises the estimate to 0.899 but implies a cost reduction of 33.9%, against 3.4% at the baseline.

The conduct parameter is strongly identified in the weak-instrument sense, with an effective first-stage F of 83.0 against the 5% critical value of 15.7 for a worst-case bias of 10% (Table S8). We also apply the conduct test of Backus et al. (2021), which compares pairs of conduct models without assuming that either is correctly specified. The test has less power here, because conduct affects markups only for the merging parties' products after the merger, which account for 277 of the 2,248 observations. With a flexible first stage, it rejects full internalization, and it never rejects the estimated $\hat{\kappa}$ (Appendix S3, Table S9). Section 7 reaches the same diagnosis at Nash pricing, at $\hat{\kappa}$ and under full internalization: at none of the three is any cap workable.

Table 5: Conduct estimates

	Single κ
Freight price	1.067***
	(0.192)
Efficiency (freight)	−0.612***
	(0.202)
κ	0.394**
	(0.193)
Demand spec	RC (origin)
κ reparametrization	Logistic
Market size	1.5× sales
Time controls	LQ trend

Notes: The Time controls row reports the controls included in the marginal-cost equation. Cost-shifter coefficients from GMM supply estimation (2-step). For logistic reparametrization (κ reparametrization = Logistic), reported κ is $\sigma(t) = e^t/(1+e^t)$ evaluated at the estimated logit-space parameter; standard errors via delta method. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6 reports average markups under observed conduct and under the Bertrand–Nash benchmark. For the RAM models, markups rise from 25.6% pre-merger to 31.4% post-merger under observed conduct, an increase of 5.8 percentage points. The remaining merged products are essentially flat, moving from 34.5% to 34.8%, while markups of non-merging

firms decline from 32.0% to 26.2% over the same period. The Nash benchmark isolates how much of the merged firms' position is conduct rather than cost or demand. Absent the increase in profit internalization, post-merger markups would have been 24.6% for RAM and 29.9% for the other merged products, so the conduct shift accounts for 6.8 and 4.9 percentage points respectively. Non-merging firms sit at 26.2% under both. Their decline reflects cost and demand conditions, not the merger.

Table 6: Marginal Cost and Markups (Pre and Post)

	Pre		Post	
	Actual	Nash	Actual	Nash
Markup (Merging, RAM)	25.588	25.588	31.409	24.600
Markup (Merging, Others)	34.469	34.469	34.837	29.947
Markup (Non Merging)	32.049	32.049	26.239	26.200
MC (Merging, RAM)	7360.010		8088.940	
MC (Merging, Others)	6415.321		7781.219	
MC (Non Merging)	4943.500		6359.399	

Notes: Markups are percentages of price, $(p-c)/p \times 100$; marginal costs in thousands of CLP. "Actual" uses the estimated conduct parameter ($\hat{\kappa} = 0.394$) post-merger and $\kappa = 0$ pre-merger. "Nash" sets $\kappa = 0$ throughout, holding marginal costs fixed. Averages over January 2015–December 2020 and January 2021–March 2023.

Table 6 also reports average estimated marginal costs. Marginal costs increase in the post-merger period for all groups of firms, with the largest absolute increase observed among non-merging firms. This generalized upward shift reflects common cost shocks affecting the industry, including higher international input prices and freight costs during the post-2020 period. The smaller increase in marginal cost for the merging parties is the estimated merger-specific efficiency, which enters through freight discounts. Relative to the no-merger counterfactual, the merger reduced marginal costs even as the common cost shock raised them in absolute terms. The merger's price increase therefore comes from the shift in conduct. The cost saving works in the opposite direction and offsets part of it.

Taken together, the point estimates indicate partial profit internalization between FCA and PSA ($\hat{\kappa} = 0.394$), with efficiency gains offsetting only part of the resulting unilateral pricing incentives. Section 7 uses these estimates to quantify the merger's welfare effects and evaluate

the remedy's structural feasibility.

7 Counterfactual Analysis

Section 3 reduces the remedy to thresholds on the capacity axis. Once the entrant is strong, so that some cap is approvable at all, the outcome turns on two questions. Can the merging parties avoid entry? And does a workable cap exist, one that clears the authority's standard, keeps the merger viable and induces entry? Those are the two coordinates of Proposition 3.1, and their answers select one of its four cases. This section establishes that the entrant is strong, answers both questions at the structural estimates, locates Chile among the cases, and measures what the realized outcome cost.

7.1 Why Toyota Did Not Enter

Counterfactual specification. To assess what Toyota's entry would have delivered, we simulate the post-merger market with Toyota as an additional firm, entering with and without a cap on its sales. Because Toyota is not in the observed sample, these counterfactuals require assumptions on its primitives. We assume it markets a single product with the observed Proace City characteristics, demand shock ξ_F imputed from the Berlingo and Partner residuals and marginal cost from the same two models' freight cost and unobserved component $\hat{\omega}$. The Proace City is built on the same platform at the same PSA facility (Section 2), so ω_F comes from the plant that would have supplied it. Toyota receives no merger-specific efficiencies. The merged entity prices at $\hat{\kappa}$ and non-merging firms play Bertrand–Nash in every scenario. Appendix C gives the full specification and Appendix S4 varies it.

Is the entrant strong? Panel A of Table 7 confirms the strong-entrant classification, with $\Delta_{\text{class}} = 10.63\%$ and a bootstrap CI of [8.08%, 12.70%] entirely above zero. $\text{CS}_{\text{unc}} > \text{CS}_0$ in every one of the 189 retained draws, conduct and Toyota's primitives held fixed.[33] Toyota's unconstrained volume is 1,869 units in 2021 and 2,188 in 2022, since the counterfactual imposes the ceiling in each calendar year and prorates it to the three months of 2023 in the sample. Both exceed the proposed cap of $Q^* = 1{,}500$ units, so the cap binds in 189 of the 189 retained draws.

[33] The intervals come from a parametric bootstrap over the demand parameters, with marginal costs and the cost-side coefficients re-estimated at each draw and $\hat{\kappa}$ and Toyota's primitives held fixed. Of 200 draws, 11 imply a negative merger efficiency, outside the range on which the thresholds are defined, and are discarded. Appendix D gives the procedure and the screening of draws.

The diagnosis is robust to Toyota's assumed primitives and to the estimation specification, both of which Appendix S4 varies. Imputing Toyota's primitives from the Partner alone, the Berlingo alone, or with merger efficiencies applied moves the point estimate of Δ_{class} between 7.32% and 14.55%, against 10.63% at the baseline. The strong-entrant classification and the sign of the remedy's welfare effect hold across both. Toyota is therefore the type the remedy is designed to use, and an approvable cap exists.

The corridor at the estimates. The lowest approvable cap is $Q_A = 620$ and the highest merger-viable cap is $Q_{MP} = 907$, so $\mathcal{C} = [Q_A, Q_{MP}]$ is an interval and the merger does not fail outright. Assumptions 3.1 and 3.2 hold at the estimates, and so does profitability without entry at the estimated conduct. The merger is profitable without entry ($\Pi^M_{\text{m,ne}} = 61{,}756.9 > \Pi_0 = 53{,}281.0$) and unprofitable under uncapped entry ($\Pi^M(Q_{\text{unc}}) = 45{,}128.5$), so a cap is necessary for viability, and it harms consumers absent a remedy, at the levels reported in (8) below.

Table 7: Classification and welfare implications

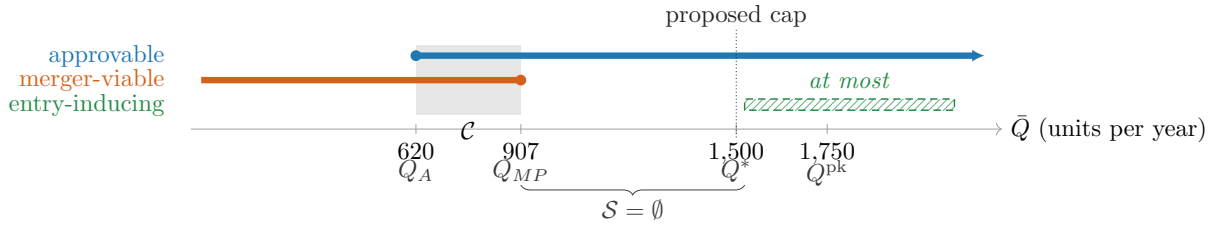
Object	Estimate	95% bootstrap CI
Panel A: Classification		
$\Delta_{\text{class}} = (CS_{unc} - CS_0)/CS_0$, strong entrant iff > 0	10.63%	[8.08%, 12.70%]
Q_{unc}, Toyota's unconstrained volume (units per year)	1,869	[1,609, 2,174]
π^F_{unc}, Toyota's operating profit under a slack cap	10,503.4	[9,008.7, 12,913.4]
$\pi^F(Q^*, E)$, Toyota's operating profit at the cap (floor on F_T)	10,509.0	[9,066.0, 12,771.9]
$\pi^F_{\max}$, Toyota's peak operating profit over caps	10,646.7	[9,168.2, 13,102.8]
Q^{pk}, cap at the peak (units)	1,750	[1,500, 2,102]
$\pi^F_{\max} - \pi^F(Q^*, E)$, width of the window for F_T	137.7	[0.0, 636.4]
Q_A, lowest approvable cap (units)	620	[502, 820]
Q_{MP}, highest merger-viable cap (units)	907	[590, 1,073]
Panel B: Welfare implications		
$\Delta_{\text{remedy}} = (CS_{m,ne} - CS_0)/CS_0$	−5.51%	[−6.67%, −4.29%]
$\Delta\Pi^M = (\Pi_{m,ne} - \Pi_0)/\Pi_0$	15.91%	[8.22%, 23.31%]

Notes: Point estimates from the non-bootstrap run. 95% confidence intervals in brackets from the 2.5/97.5 percentiles of the parametric bootstrap (189 valid draws). $\pi^F_{\max} = 10{,}646.7$. Section 7 discusses what the interval $(\pi^F(Q^*, E), \pi^F_{\max}]$ does and does not identify.

Can the parties avoid entry? They can. Entry is unavoidable only when every approvable capacity induces it. Observed non-entry at the proposed cap places F_T above $\pi^F(Q^*, E) =$ 10,509.0, and the entrant's profit is still rising there, so it earns less at every approvable capacity up to Q^*, the lowest one included ($\pi^F(Q_A, E) = 7{,}358.8$, below $\pi^F(Q^*, E)$ in every draw). Those capacities deter, so $\mathcal{D}$, the approvable capacities that deter entry, is not empty and the parties are free to propose one of them. This is the first coordinate of Proposition 3.1, and it already rules out cases (i) and (ii).

Does a workable cap exist? None does. Observed non-entry at a binding cap reveals the one-sided bound $F_T > \pi^F(Q^*, E) = 10{,}509.0$. The entrant's operating profit anywhere inside $[Q_A, Q_{MP}]$ is at most $\pi^F(Q_{MP}, E) = 8{,}920.2$, below that floor in every draw, so the corridor and the entry-inducing set are disjoint, $\mathcal{S} = \emptyset$ (Figure 3). Both profits in that comparison accumulate over the same window, so the disjointness does not depend on the horizon, though the level of the bound does. Toyota gave up at least the operating profit the 27-month figures record, and more if the remedy's term ran past the sample. The corridor sits on the upward arm of the entrant's profit, which is still increasing across it (7,358.8 at Q_A, 8,920.2 at Q_{MP}) and peaks only at 1,750 units, above $Q_{MP} = 907$, at 10,646.7 (Figure S2). Tightening the cap to preserve viability moves the entrant further from participation, not closer. The remedy's failure is therefore not the choice of cap. No cap the authority could approve would have drawn Toyota in while keeping the merger viable. Some cap outside the corridor could have induced entry only if F_T lies between 10,509.0 and 10,646.7, which non-entry at 1,500 does not rule out (Proposition A.2(iv)). Any such cap lies above Q_{MP} and would have left the merger unprofitable. Figure 4 in Appendix G plots the underlying curves with bootstrap bands.

Figure 3: The three constraint sets at the structural estimates



Notes: the empirical counterpart of panel (d) of Figure 1, drawn to scale on the capacity axis. Light shading is the corridor $\mathcal{C} = [Q_A, Q_{MP}]$. The entry-inducing set is hatched because F_T is only bounded below; it is drawn at its widest, and is empty if F_T exceeds π^F's peak.

The classification. With entry avoidable and $\mathcal{S} = \emptyset$, the FCA–PSA merger in Chile is case (iv) of Proposition 3.1, structural failure, a strong entrant for which no workable cap exists.[34] The observed actions are the ones the case predicts. The parties merged, proposed an approvable capacity that did not induce entry, and Toyota stayed out, so realized surplus is $\mathrm{CS}_{\mathrm{m,ne}}$. The proposed cap of 1,500 lies above $Q_{MP} = 907$, so had Toyota entered the merger would not have remained viable. The parties would have been left 8.37% below the no-merger benchmark, and the cap exceeds Q_{MP} in 189 of the 189 retained draws. Non-entry at any operative cap between 907 and 1,750 units places Q_F above it on the increasing branch of π^F, and $\mathcal{S} = \emptyset$ follows.

What the merger is compared against. Viability is measured against $\Pi_0 = 53{,}281.0$, the profit the merging parties earn without the merger. The transaction could not close without Chile's approval (Section 2), so the alternative to an approved cap was no merger. Two other benchmarks lie on either side of Π_0. If the merger proceeded abroad and the two lines were still priced independently in Chile, the freight saving would reach both lines at $\kappa = 0$ and the parties would earn 55,470.5, 4.1% above Π_0. The highest merger-viable cap then falls from 907 to 652 units, and no workable cap exists, by a wider margin.[35] If instead viability in Chile did not constrain the parties, they would accept any cap, and the diagnosis turns on where F_T lies. Above Toyota's peak operating profit of 10,646.7, no cap induces entry and the case remains structural failure. In the window of 137.7 between that peak and the floor that non-entry at 1,500 establishes, a workable cap exists, the case is strategic deterrence, and conditioning approval on realized entry would have brought Toyota in. The data do not say which. What does not depend on it is the outcome. Under either reading the parties could deter entry, consumers received $\mathrm{CS}_{\mathrm{m,ne}}$, and conditioning approval on entry would have left them at least at CS_0, by drawing Toyota in or by preventing the merger.

7.2 What It Cost

We now measure what Toyota's absence cost consumers by comparing consumer surplus across the four market configurations. Profits and surplus are accumulated over the post-

[34] Under the nested logit on origin (Table 4) the merger's harm is similar (−6.11% against −5.51%) but entry adds little, so $\Delta_{\mathrm{class}} = -4.19\%$ and the entrant is weak: the two models differ on how much consumers value a product that duplicates the Berlingo, not on the merger. $\mathcal{S}$ is again empty, now because no capacity is approvable at all, and the remedy fails because entry could not have restored surplus rather than because no workable capacity existed. The specification we report rationalizes all three observed actions, approval, merger and non-entry, as case (iv); the nested logit rationalizes only the last.

[35] Consumers in that configuration are 2.9% above the no-merger level (Table 12). The authority's standard is still CS_0, so Q_A does not move.

merger window, from January 2021 to March 2023, and expressed in millions of CLP (Table 14 in Appendix G). The authority evaluated the remedy assuming entry at the cap, which yields $\mathrm{CS}(Q^*, E) = 145{,}268$, whereas the realized outcome without entry delivered $\mathrm{CS}_{\mathrm{m,ne}} = 128{,}580$. The four levels are ordered as follows:

$$\underbrace{\mathrm{CS}_{\mathrm{unc}}}_{150{,}543} > \underbrace{\mathrm{CS}(Q^*, E)}_{145{,}268} > \underbrace{\mathrm{CS}_0}_{136{,}078} > \underbrace{\mathrm{CS}_{\mathrm{m,ne}}}_{128{,}580}. \tag{8}$$

The top two are continuations that never occurred: uncapped entry, which would have made the merger unprofitable, and entry at the cap, which the authority credited. Its evaluation was correct in isolation, $\mathrm{CS}(Q^*, E) > \mathrm{CS}_0$, but conditional on an event that did not happen. Three of the four alternatives in Proposition 3.2 would have left consumers at $\mathrm{CS}_0 = 136{,}078$. Without a remedy the merger fails the screen, and verified approval and an authority-set capacity both prevent it here, because no workable cap exists. Uncapped access would not have helped either. It delivers $\mathrm{CS}_{\mathrm{m,ne}}$ or CS_0 and never more (Proposition 3.2, regime (c)).

Conduct and the harm. Holding the market configuration fixed and varying only conduct (Table 12), consumer surplus stands at +2.9% under Nash pricing with efficiencies ($\kappa = 0$), −5.51% at $\hat{\kappa}$, and −17.6% under full internalization, and it is negative for κ above 0.14. The estimate $\hat{\kappa} = 0.394$ (0.193) rejects both Bertrand–Nash and full internalization at the 5% level and lies inside that region, as the authority's decision to require a remedy presumes. Efficiencies offset roughly a third of the unilateral effect (Table 15). Total surplus nonetheless rises by 0.9%, since the cost saving (2.4%) outweighs the deadweight loss from the conduct shift (−1.5%); under the FNE's consumer-surplus standard that does not alter the assessment.

Across degrees of internalization. The corridor moves with conduct as well (Table 13). At Nash conduct the merger clears the screen on its own, so no remedy is called for. At the estimated conduct the corridor is [620, 907] and contains no entry-inducing cap. At full internalization the harm is large enough that approval requires a cap of at least 1,694 units while viability allows at most 1,422, so no cap is both approvable and merger-viable and the merger fails outright. Toyota's entry profit rises in κ (Table 12), so high internalization is the case least favorable to the absence of a workable cap. At none of the three degrees of internalization is a cap workable, and at none did the remedy deliver anything to consumers.

8 Conclusion

Entry-facilitating remedies allow competition authorities to clear mergers that raise competitive concerns when divestitures are not feasible or proportionate, by giving a new competitor the means to enter. This paper asks whether such remedies protect consumers when, as in many jurisdictions, the merging parties set the terms of access and the authority approves the merger before the entrant decides whether to enter.

Our model shows that a strong entrant is necessary for the remedy to be approvable, but what decides whether consumers are protected is whether the merging parties can avoid its entry. The parties choose how much it may sell. When they can propose a cap that the authority accepts and the entrant declines, they do so, and the merger goes through without entry. When they cannot, the remedy works, but it restores consumers only to their position without the merger, since the parties concede no more than the authority requires.

The FCA–PSA merger in Chile illustrates the first case. It harmed consumers. The harm does not appear in a reduced-form comparison of the parties' prices with those of non-merging vans, which nets the increase in market power against the parties' cost savings. The structural model separates the two: the merged firm partially internalized the profits of its partner's models, and cost savings offset only about a third of the resulting price increase. Toyota was a strong entrant, as its unrestricted entry would have left consumers 10.63% better off than without the merger, yet it did not enter, and consumers ended up 5.51% worse off. Its absence reflects a structural failure. At any cap that kept the merger profitable in Chile, Toyota would have earned less than at the cap it declined, so no cap could have satisfied the authority, kept the merger profitable and attracted Toyota. The remedy's effect was to clear a merger that the authority's initial assessment would have blocked.

Two limits qualify this diagnosis. First, it rests on a single decision by a single firm. Toyota did not enter during the more than two years of the remedy's term covered by our sample, and it had still not entered more than five years after approval. That decision could also reflect reasons specific to Toyota that are unrelated to the design of the remedy. Second, the lower bound on Toyota's entry cost comes from the structural model rather than from an independent estimate. Neither the harm nor Toyota's classification as a strong entrant depends on that bound. The absence of a workable cap does, through a threshold: it requires an entry cost above Toyota's profit at the largest cap the merging parties would accept, and the bound exceeds that profit in every draw (Section 7.1).

Two changes to the procedure would have avoided the outcome in Chile. An authority-set cap requires the authority to know, before the merger is implemented, Toyota's profit under each possible cap, its cost of entering and the parties' profits once it entered. Conditioning approval on realized entry requires none of that information, only approval that can be withheld when entry fails and access that is monitored. A cap the entrant declines then no longer secures approval, so consumers are never left worse off than without the merger. In Chile, under either change, consumers would have kept their no-merger surplus, which at our estimates means the parties would not have merged.

The remedy studied here is specific, but its structure is not. Divestitures to buyers not committed in advance, licensing obligations whose royalty and field of use the parties draft, and supply agreements whose volume and term they set all depend on a third party who must choose to take them up, on terms the merging parties write, and are approved as though that choice had already been made. In each case, the terms must satisfy the authority, keep the merger worth doing and make the third party willing to take them up. Competition authorities already treat the last condition as the fragile one in divestitures, where upfront-buyer requirements test it before approval (European Commission, 2008). Entry-facilitating remedies carry no equivalent safeguard, even though the entrant's willingness is the only condition the authority cannot observe when it decides, and the one the party drafting the terms gains from failing.

References

Ashenfelter, Orley C., and Daniel S. Hosken. 2010. "The Effect of Mergers on Consumer Prices: Evidence from Five Mergers on the Enforcement Margin." *Journal of Law and Economics* 53 (3): 417–466.

Ashenfelter, Orley C., Daniel S. Hosken, and Matthew C. Weinberg. 2013. "The Price Effects of a Large Merger of Manufacturers: A Case Study of Maytag–Whirlpool." *American Economic Journal: Economic Policy* 5 (1): 239–261.

Backus, Matthew, Christopher T. Conlon, and Michael Sinkinson. 2021. "Common Ownership and Competition in the Ready-to-Eat Cereal Industry." Working Paper 28350, National Bureau of Economic Research.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–890.

Bhattacharya, Vivek, Gastón Illanes, Avner A. Kreps, José D. Salas, and David Stillerman. 2026. "A Large-Scale Evaluation of Merger Simulations." Working Paper 35473, National Bureau of Economic Research.

Björnerstedt, Jonas, and Frank Verboven. 2016. "Does Merger Simulation Work? Evidence from the Swedish Analgesics Market." *American Economic Journal: Applied Economics* 8 (3): 125–164.

Brambilla, Irene. 2005. "A Customs Union with Multinational Firms: The Automobile Market in Argentina and Brazil."

Cabral, Luís M. B. 2003. "Horizontal Mergers with Free-Entry: Why Cost Efficiencies May Be a Weak Defense and Asset Sales a Poor Remedy." *International Journal of Industrial Organization* 21 (5): 607–623.

Calder-Wang, Sophie, and Gi Heung Kim. 2024. "Algorithmic Pricing in Multifamily Rentals: Efficiency Gains or Price Coordination?", SSRN Working Paper 4403058.

Caradonna, Peter, Nathan H. Miller, and Gloria Sheu. 2025. "Mergers, Entry, and Consumer Welfare." *American Economic Journal: Microeconomics* 17 (3): 103–130.

Chamberlain, Gary. 1987. "Asymptotic Efficiency in Estimation with Conditional Moment Restrictions." *Journal of Econometrics* 34 (3): 305–334.

Chaves, Daniel. 2022. "Market Power, Taxation and Product Variety in the Brazilian Automobile Industry."

Ciliberto, Federico, and Jonathan W. Williams. 2014. "Does Multimarket Contact Facilitate Tacit Collusion? Inference on Conduct Parameters in the Airline Industry." *RAND Journal of Economics* 45 (4): 764–791.

Cisternas, Rosario. 2018. "Estimación de pass-through: caso de colusión."

Competition and Markets Authority, Bundeskartellamt, and Australian Competition and Consumer Commission. 2021. "Joint Statement on Merger Control Enforcement."

Conlon, Christopher, and Jeff Gortmaker. 2020. "Best Practices for Differentiated Products Demand Estimation with PyBLP." *RAND Journal of Economics* 51 (4): 1108–1161.

Coşar, A. Kerem, Paul L. E. Grieco, Shengyu Li, and Felix Tintelnot. 2018. "What drives home market advantage?" *Journal of International Economics* 110, 135–150.

Cosnita, Andreea, and Jean-Philippe Tropeano. 2009. "Negotiating Remedies: Revealing the Merger Efficiency Gains." *International Journal of Industrial Organization* 27 (2): 188–196.

Cosnita-Langlais, Andreea, and Jean-Philippe Tropeano. 2012. "Do Remedies Affect the Efficiency Defense? An Optimal Merger-Control Analysis." *International Journal of Industrial Organization* 30 (1): 58–66.

Delaprez, Yann. 2026. "Can Merger Remedies Cause More Harm than Good?" *Journal of Law, Economics, and Organization*, Forthcoming.

Department for Transport. 2024. "Understanding the Road Freight Market." technical report, Department for Transport.

Dertwinkel-Kalt, Markus, and Christian Wey. 2016. "Merger Remedies in Oligopoly under a Consumer Welfare Standard." *Journal of Law, Economics, and Organization* 32 (1): 150–179.

Eisenberg, Thomas, Manuel Estay, and Debi Prasad Mohapatra. 2026. "Welfare Effects of Trade Associations: The Case of the Chilean Salmon Export Industry." *RAND Journal of Economics* 57 (3): 745–763.

European Commission. 2008. "Commission Notice on Remedies Acceptable under Council Regulation (EC) No. 139/2004."

——— 2021. "Decision in Case M.9730 – FCA/PSA under Regulation (EC) No. 139/2004." Official Journal of the European Union, C 423.

Farrell, Joseph. 2003. "Negotiation and Merger Remedies: Some Problems." In *Merger Remedies in American and European Union Competition Law*, edited by François Lévêque, and Howard Shelanski Edward Elgar, Chap. 6, 95–105.

Federal Trade Commission. 2017. "The FTC's Merger Remedies 2006–2012: A Report of the Bureaus of Competition and Economics." technical report, Federal Trade Commission.

Fiat Chrysler Automobiles N.V.. 2020. "Merger of Peugeot S.A. with and into Fiat Chrysler Automobiles N.V.: Prospectus." Dated November 23, 2020. Filed with the U.S. Securities and Exchange Commission pursuant to Rule 424(b)(3), Registration No. 333-240094.

Fiscalía Nacional Económica. 2020. "Merger Approval Report with Remedies: Fiat Chrysler Automobiles N.V. and Peugeot S.A. (Case FNE F233-2020)."

Friberg, Richard, and André Romahn. 2015. "Divestiture Requirements as a Tool for Compe-

tition Policy: A Case from the Swedish Beer Market." *International Journal of Industrial Organization* 42, 1–18.

Gandhi, Amit, and Jean-François Houde. 2019. "Measuring Substitution Patterns in Differentiated-Products Industries." Working Paper 26375, National Bureau of Economic Research.

Global Competition Review. 2021. "GCR Awards 2021."

Goldberg, Pinelopi Koujianou, and Frank Verboven. 2001. "The evolution of price dispersion in the European car market." *The Review of Economic Studies* 68 (4): 811–848.

Grieco, Paul L. E., Charles Murry, and Ali Yurukoglu. 2024. "The Evolution of Market Power in the U.S. Automobile Industry." *Quarterly Journal of Economics* 139 (2): 1201–1253.

Grigolon, Laura, Mathias Reynaert, and Frank Verboven. 2018. "Consumer Valuation of Fuel Costs and Tax Policy." *American Economic Journal: Economic Policy* 10 (3): 193–225.

Grigolon, Laura, and Frank Verboven. 2014. "Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation." *Review of Economics and Statistics* 96 (5): 916–935.

Gómez-Lobo, Andrés, and J. J. Price. 2025. "Environmental Impact of a Green Tax on New Vehicle Sales in the Presence of Tax Exemptions." Working Paper.

Houde, Jean-François. 2012. "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *American Economic Review* 102 (5): 2147–2182.

International Competition Network. 2016. "Merger Remedies Guide." technical report, International Competition Network.

Ivaldi, Marc, and Frank Verboven. 2005. "Quantifying the Effects from Horizontal Mergers in European Competition Policy." *International Journal of Industrial Organization* 23 (9–10): 669–691.

Jara, Oscar, and Pedro Rojas. 2026. "The Welfare Effects of Structural Remedies." Unpublished manuscript, Norwegian School of Economics and Toulouse School of Economics.

Lagos, Vicente. 2018. "Effectiveness of Merger Remedies: Evidence from the Retail Gasoline Industry." *Journal of Industrial Economics* 66 (4): 942–980.

Lyons, Bruce, and Andrei Medvedev. 2007. "Bargaining over Remedies in Merger Regulation." CCP Working Paper 07-3, ESRC Centre for Competition Policy, University of East Anglia.

Maier-Rigaud, Frank, and Benedikt Loertscher. 2020. "Structural vs. Behavioral Remedies." *CPI Antitrust Chronicle*.

Michel, Christian. 2016. "Identification and Estimation of Intra-Firm and Industry Competition via Ownership Change." Working Paper.

Michel, Christian, José M. Paz y Miño, and Stefan Weiergraeber. 2024. "Estimating Industry Conduct Using Promotion Data." *RAND Journal of Economics* 55 (4): 658–683.

Miller, Nathan H., and Matthew C. Weinberg. 2017. "Understanding the Price Effects of the MillerCoors Joint Venture." *Econometrica* 85 (6): 1763–1791.

Montag, Felix. 2023. "Mergers, Foreign Competition, and Jobs: Evidence from the U.S. Appliance Industry." Working Paper 326, Stigler Center.

Montiel Olea, José Luis, and Carolin Pflueger. 2013. "A Robust Test for Weak Instruments." *Journal of Business & Economic Statistics* 31 (3): 358–369.

Morrow, W. Ross, and Steven J. Skerlos. 2011. "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand." *Operations Research* 59 (2): 328–345.

Motta, Massimo, Michele Polo, and Helder Vasconcelos. 2002. "Merger Remedies in the European Union: An Overview." Conference paper.

Nevo, Aviv. 2000. "Mergers with Differentiated Products: The Case of the Ready-to-Eat Cereal Industry." *RAND Journal of Economics* 31 (3): 395–421.

Nocke, Volker, and Andrew Rhodes. 2026. "Optimal Merger Remedies." TSE Working Paper 25-1660, Toulouse School of Economics, Also CEPR Discussion Paper 20573.

Nocke, Volker, and Nicolas Schutz. 2018. "Multiproduct-firm oligopoly: An aggregative games approach." *Econometrica* 86 (2): 523–557.

——— 2025. "An Aggregative Games Approach to Merger Analysis in Multiproduct-Firm Oligopoly." *The RAND Journal of Economics* 56 (3): 233–250.

Ohashi, Hiroshi, and Yusuke Toyama. 2017. "The Effects of Domestic Merger on Exports: A Case Study of the 1998 Korean Automobile Industry." *Journal of International Economics* 107, 147–164.

Organisation for Economic Co-operation and Development. 2012. "Remedies in Merger Cases:

Key Findings, Summary and Notes." Roundtables on Competition Policy Papers 125, OECD.

Osinski, Felipe D., and Joshua Sandford. 2020. "Evaluating Mergers and Divestitures: A Casino Case Study." SSRN Electronic Journal.

Reynaert, Mathias, and Frank Verboven. 2014. "Improving the Performance of Random Coefficients Demand Models: The Role of Optimal Instruments." *Journal of Econometrics* 179 (1): 83–98.

Stellantis. 2021. "The Merger of FCA and Groupe PSA Has Been Completed." Press release.

——— 2022. "Stellantis and Toyota Expand Partnership with New Large-Size Commercial Van, Including an Electric Version." Press release, May 30, 2022.

Taylor, Christopher T., and Daniel S. Hosken. 2007. "The Economic Effects of the Marathon–Ashland Joint Venture." *Journal of Industrial Economics* 55 (3): 419–451.

U.S. Department of Justice, Antitrust Division. 2020. "Merger Remedies Manual." technical report, U.S. Department of Justice.

Vasconcelos, Helder. 2010. "Efficiency Gains and Structural Remedies in Merger Control." *Journal of Industrial Economics* 58 (4): 742–766.

Verboven, Frank. 1996. "International price discrimination in the European car market." *The RAND Journal of Economics*, 240–268.

Weinberg, Matthew C., and Daniel S. Hosken. 2013. "Evidence on the Accuracy of Merger Simulations." *Review of Economics and Statistics* 95 (5): 1584–1600.

Werden, Gregory J., and Luke M. Froeb. 1998. "The Entry-Inducing Effects of Horizontal Mergers: An Exploratory Analysis." *Journal of Industrial Economics* 46 (4): 525–543.

Wollmann, Thomas G. 2018. "Trucks without bailouts: Equilibrium product characteristics for commercial vehicles." *American Economic Review* 108 (6): 1364–1406.

Appendix

A Bertrand Subgame and Supporting Results

A.1 Unconstrained Bertrand Equilibrium

Write $T^f \equiv \sum_{j\in\mathcal{J}_f} \exp(v_j - \alpha c_j)$ for firm f's *type*, the quality-adjusted cost aggregate of its products, where v_j is mean utility net of price and c_j marginal cost. The merger lowers the marginal cost of each of the parties' products from c_j to $c_j - E$, so the merged entity's type is

$$T^M(E) \;\equiv\; \sum_{j\in\mathcal{J}_1\cup\mathcal{J}_2} \exp\Big(v_j - \alpha(c_j - E)\Big) \;=\; e^{\alpha E}\,(T^1 + T^2),$$

with $T^M(0) = T^1 + T^2$, the merger without efficiencies, and T^M strictly increasing in E. A product-specific saving, as in the estimated cost function where it varies with freight cost, enters each term as a positive multiplier and raises T^M in the same way; the results below use only this monotonicity. T^F is the entrant's type. Under MNL demand each firm f sets price to satisfy the markup condition $\mu^f = m(T^f/H)$, where $m(x)$ is the unique positive solution to:

$$\mu\Big(1 - xe^{-\mu}\Big) \;=\; 1, \tag{9}$$

and $S(x) \equiv xe^{-m(x)}$ is the equilibrium market share. The aggregator H satisfies the adding-up constraint:

$$\frac{1}{H} + \sum_{f\in\{M\}\cup\mathcal{F}_0} S\!\left(\frac{T^f}{H}\right) \;=\; 1. \tag{10}$$

Since m is strictly increasing and $S(x) = 1 - 1/m(x)$, each $S(T^f/H)$ is strictly decreasing in H. The left-hand side of (10) therefore falls strictly from $+\infty$ to 0, so the constraint has a unique positive root.

Consumer surplus is $\mathrm{CS}(H) = \frac{1}{\alpha}\ln H$. The merged entity's profit is:

$$\Pi^M(E) \;=\; \frac{1}{\alpha}\left(m\!\left(\frac{T^M}{H}\right) - 1\right).$$

The pre-merger aggregator H_0 and the unconstrained post-merger aggregator H_{unc} each solve

(10) under their respective type vectors: H_0 uses firms 1 and 2 as separate entities (replacing $\{M\}$ with $\{1,2\}$), while H_{unc} additionally includes Toyota (adding $S(T^F/H)$ to the left-hand side). The capped-entry aggregator $H(\bar{Q}, E)$ solves the modified constraint (AUP) defined in Section A.2 below.

A.2 Equilibrium under a Binding Cap

When $\bar{Q} < Q_{\text{unc}}$, the cap binds. Toyota's price satisfies the binding pricing condition, in which v^F is the entrant's mean utility net of price:

$$\exp(v^F - \alpha p^F) \;=\; \bar{Q} \cdot H, \tag{BP}$$

yielding Toyota's markup $\mu^F_{\text{cap}} = \ln(T^F/(\bar{Q}H))$. Quantities in this appendix are shares of market size, so the cap $\bar{Q}$ enters the aggregator directly. The unit figures reported elsewhere are the corresponding annual volumes. The capped adding-up constraint is:

$$\frac{1}{H} + \bar{Q} + \sum_{f \in \{M\} \cup \mathcal{F}_0} S\left(\frac{T^f}{H}\right) \;=\; 1. \tag{AUP}$$

The cap replaces Toyota's share $S(T^F/H)$ with the constant $\bar{Q}$, so the left-hand side remains strictly decreasing, now falling from $+\infty$ to $\bar{Q} < 1$. Hence (AUP) has a unique positive root $H(\bar{Q}, E)$. At $\bar{Q} = 0$, (AUP) reduces to (10), whose root is the merger-without-entry aggregator, so $\text{CS}(0, E) = \text{CS}_{\text{m,ne}}$ and $\Pi^M(0, E) = \Pi^M_{\text{m,ne}}$.

Toyota's operating profit under the binding cap is:

$$\pi^F(\bar{Q}, E) \;=\; \frac{\bar{Q}}{\alpha} \ln\left(\frac{T^F}{\bar{Q}\, H(\bar{Q}, E)}\right), \tag{11}$$

and the merged entity's profit is:

$$\Pi^M(\bar{Q}, E) \;=\; \frac{1}{\alpha}\left(m\left(\frac{T^M}{H(\bar{Q}, E)}\right) - 1\right). \tag{12}$$

When $\bar{Q} \geq Q_{\text{unc}}$ the ceiling is slack. Toyota's unconstrained best response sells $Q_{\text{unc}} \leq \bar{Q}$, the cap does not bind, and the equilibrium coincides with the unconstrained Bertrand equilibrium of Section A.1. Hence $H(\bar{Q}, E) = H_{\text{unc}}$, $\pi^F(\bar{Q}, E) = \pi^F_{\text{unc}}$, and $\Pi^M(\bar{Q}, E) = \Pi^M(Q_{\text{unc}}, E)$ are all constant in $\bar{Q}$ on the slack region. At the boundary $\bar{Q} = Q_{\text{unc}}$ the binding pricing

condition (BP) reduces to Toyota's unconstrained markup $\mu^F = m(T^F/H_{\text{unc}})$, so the binding and unconstrained aggregators coincide and H, π^F, and Π^M are continuous there. The transition from the binding to the slack regime is therefore a kink in these objects, not a jump.

A.3 Pre-Merger Baseline and Monotonicity

The pre-merger aggregator H_0 solves (10) with firms 1, 2 separate (no merger, no Toyota). Consumer surplus is $\text{CS}_0 = \frac{1}{\alpha} \ln H_0$ and pre-merger profits are $\Pi^i_{\text{pre}} = \frac{1}{\alpha}(m(T^i/H_0) - 1)$ for $i \in \{1, 2\}$.

The three monotonicity properties follow from Nocke and Schutz (2018), Proposition 6 ($m' > 0$ and $S' > 0$), applied to (10) and (AUP) via the implicit function theorem:

(M1) $\partial H(\bar{Q}, E)/\partial \bar{Q} > 0$, equivalently $\partial\, \text{CS}(\bar{Q}, E)/\partial \bar{Q} > 0$: a looser cap raises H and hence consumer surplus.

(M2) $\partial H(\bar{Q}, E)/\partial E > 0$: higher efficiency raises H through a higher T^M.

(M3) $\partial \Pi^M(\bar{Q}, E)/\partial \bar{Q} < 0$ and $\partial \Pi^M(\bar{Q}, E)/\partial E > 0$: the merging parties prefer a tighter cap and benefit from higher efficiency.

These derivatives are taken on the binding region $\bar{Q} < Q_{\text{unc}}$. On the slack region $\bar{Q} \geq Q_{\text{unc}}$ the objects are constant in $\bar{Q}$ (Section A.2), so H and CS are non-decreasing and Π^M non-increasing in $\bar{Q}$ globally, each continuous with a kink at Q_{unc}. In particular Π^M attains its minimum over caps at Q_{unc}, so any profit-neutrality cap satisfies $Q_{MP} \leq Q_{\text{unc}}$ and no merger-viable threshold lies in the slack region.

A.4 Parameter Space Characterization

Lemma A.1 (Efficiency and Regime Classification)**.** (i) There exists a unique $\tilde{E} > 0$ with

$$\text{CS}_{\text{m,ne}}(\tilde{E}) = \text{CS}_0, \tag{13}$$

such that for $E \geq \tilde{E}$ the merger passes the unconditional screen without a remedy. (ii) If Assumption 3.2 holds at some $E \geq 0$, there exists a unique $\bar{E} > 0$ with

$$\Pi^M(Q_{\text{unc}}, \bar{E}) = \Pi_0, \tag{14}$$

such that for $E \geq \bar{E}$ the merger is profitable under uncapped Toyota entry and a cap is unnecessary. The admissible region is $\mathcal{A} \equiv \{E : E < E^*\}$, where $E^* \equiv \min(\tilde{E}, \bar{E})$.

Proof. *(i)* $\mathrm{CS}_{\mathrm{m,ne}}$ is strictly increasing in E (M2) and unbounded above; at $E = 0$, $\mathrm{CS}_{\mathrm{m,ne}}(0) < \mathrm{CS}_0$, since a merger without efficiencies lowers consumer surplus (Nocke and Schutz, 2025, Proposition 2). By the intermediate value theorem (IVT) and strict monotonicity, $\tilde{E}$ exists and is unique.

(ii) $\Pi^M(Q_{\mathrm{unc}}, E)$, the merged entity's profit under uncapped entry, is strictly increasing in E, by the argument of (M3) applied to (10) with the entrant included, and unbounded above. At $E = 0$ it lies below Π_0 whenever uncapped entry restores consumer surplus without efficiencies, by CMS Proposition 5, and in general whenever Assumption 3.2 holds at some $E \geq 0$, by monotonicity. The IVT gives a unique $\bar{E}$. □

A.5 Toyota's Variable Operating Profit

Proposition A.2 (Toyota's Variable Operating Profit)**.** $\pi^F(\bar{Q}, E)$ satisfies: (i) it is strictly increasing in $\bar{Q}$ for $\bar{Q}$ small and strictly decreasing as $\bar{Q} \uparrow Q_{\mathrm{unc}}$, and constant at π^F_{unc} for $\bar{Q} \geq Q_{\mathrm{unc}}$; hence it attains an interior maximum at some $Q^{\mathrm{pk}}(E) < Q_{\mathrm{unc}}$; (ii) strictly decreasing in E; (iii) $\pi^F_{\max} \equiv \pi^F(Q^{\mathrm{pk}}, E) > \pi^F_{\mathrm{unc}}$, so the entrant's operating profit is maximized by a cap that binds strictly and a slack cap does not deliver its highest attainable profit; (iv) if $\pi^F_{\max} \geq F_T$ and π^F is single-peaked in $\bar{Q}$ on $[0, Q_{\mathrm{unc}}]$, strictly increasing below Q^{pk} and strictly decreasing above it, the set of caps satisfying (F) is the interval $[Q_F(E), Q_F^+(E)]$ containing Q^{pk}, where $\pi^F(Q_F, E) = \pi^F(Q_F^+, E) = F_T$ and $Q_F^+ = +\infty$ whenever $F_T \leq \pi^F_{\mathrm{unc}}$; Q_F is increasing and Q_F^+ decreasing in E. If $\pi^F_{\max} < F_T$, (F) fails at every cap.

Proof. *(i)* $\pi^F(\cdot, E)$ is continuous on $[0, Q_{\mathrm{unc}}]$, with $\pi^F(0, E) = 0$ since $\bar{Q}\ln(1/\bar{Q}) \to 0$. Differentiating (11) with respect to $\bar{Q}$ holding H fixed gives the direct effect $(\mu^F_{\mathrm{cap}} - 1)/\alpha > 0$, which dominates for $\bar{Q}$ small. As the cap ceases to bind the direct effect does not vanish, since by (9) $\mu^F_{\mathrm{cap}} \to 1/(1 - Q_{\mathrm{unc}}) > 1$, but it is dominated by the indirect effect through the aggregator H, which by (M1) a looser cap raises, with $\partial\pi^F/\partial H = -\bar{Q}/(\alpha H) < 0$. Implicit differentiation of the capped adding-up constraint (AUP) gives $H'(\bar{Q}) = H/D$ with $D = 1/H + \sum_{f \neq F} x_f S'(x_f)$, and since $xS'(x) = s(1-s)^2/((1-s)^2 + s) < s$ we have $D < 1 - Q_{\mathrm{unc}}$. Evaluating at $\bar{Q} \to Q^-_{\mathrm{unc}}$ gives $d\pi^F/d\bar{Q} = (Q_{\mathrm{unc}}/\alpha)\,[1/(1 - Q_{\mathrm{unc}}) - 1/D] < 0$, so $d\pi^F/d\bar{Q} < 0$ in a neighborhood of Q_{unc}. The derivative therefore changes sign on $(0, Q_{\mathrm{unc}})$, and the maximum of π^F on $[0, Q_{\mathrm{unc}}]$, which exists by continuity, is interior. *(ii)* On a binding cap, higher E raises H by (M2),

and $\partial\pi^F/\partial H < 0$. On the slack region $\pi^F = \pi^F_{\text{unc}} = (m(T^F/H_{\text{unc}}) - 1)/\alpha$, which falls in E because $m' > 0$ and H_{unc} rises with T^M by the same argument applied to (10). *(iii)* Immediate from (i), since $\pi^F(\bar{Q}, E) = \pi^F_{\text{unc}}$ for all $\bar{Q} \geq Q_{\text{unc}}$ and π^F is strictly decreasing on a left neighborhood of Q_{unc}. *(iv)* Follows from (iii) and the IVT on each branch, starting from $\pi^F(0, E) = 0 < F_T$, with comparative statics in E from (ii). □

The single-peakedness hypothesis of part (iv) is verified numerically at the estimates rather than proved.

B Full Equilibrium Characterization

Throughout this appendix Assumptions 3.1 and 3.2 hold and $\pi^F(\cdot, E)$ is single-peaked on $[0, Q_{\text{unc}}]$, the hypothesis of Proposition A.2(iv), under which the entry-inducing set is the interval $[Q_F, Q_F^+]$.

B.1 Corridor Endpoints and Efficiency

The level sets $\text{CS}(\bar{Q}, E) = \text{CS}_0$ and $\Pi^M(\bar{Q}, E) = \Pi_0$ locate the corridor endpoints of Section 3.3 in $(\bar{Q}, E)$ space. Read as functions of E, they give the dependence of Q_A and Q_{MP} on efficiency used there.

Lemma B.1 (Corridor Endpoints and Efficiency)**.** (i) For every $E \in [0, \tilde{E}]$ at which the entrant is strong, $\text{CS}_{\text{unc}}(E) \geq \text{CS}_0$, there is a unique lowest approvable cap $Q_A(E)$ with $\text{CS}(Q_A, E) = \text{CS}_0$; it is strictly decreasing in E, with $Q_A(\tilde{E}) = 0$. (ii) For every admissible E there is a unique highest merger-viable cap $Q_{MP}(E)$ with $\Pi^M(Q_{MP}, E) = \Pi_0$; it is strictly increasing in E. A more efficient merger therefore clears the standard at a smaller cap and stays viable at a larger one.

Proof. (i) On the binding region $\text{CS}(\cdot, E)$ rises continuously from $\text{CS}_{\text{m,ne}}(E) \leq \text{CS}_0$, as $E \leq \tilde{E}$, to $\text{CS}_{\text{unc}}(E) \geq \text{CS}_0$, as the entrant is strong, so by (M1) and the IVT it crosses CS_0 exactly once. Since $\text{CS} = \frac{1}{\alpha} \ln H$ is strictly increasing in $\bar{Q}$ (M1) and in E (M2) on the binding region, the implicit function theorem gives $dQ_A/dE = -(\partial\,\text{CS}/\partial E)/(\partial\,\text{CS}/\partial\bar{Q}) < 0$, and $Q_A(\tilde{E}) = 0$ is the definition of $\tilde{E}$ (Lemma A.1).

(ii) On the binding region $\Pi^M(\cdot, E)$ falls continuously by (M3) from $\Pi^M(0, E) = \Pi^M_{\text{m,ne}} > \Pi_0$ to $\Pi^M(Q_{\text{unc}}, E) < \Pi_0$, as E is admissible, so it crosses Π_0 exactly once. Since Π^M is strictly

decreasing in $\bar{Q}$ and strictly increasing in E on the binding region (M3), the implicit function theorem gives $dQ_{MP}/dE = -(\partial\Pi^M/\partial E)/(\partial\Pi^M/\partial\bar{Q}) > 0$. □

B.2 Weak Entrant

Proposition B.2 (Merger Abandoned: Weak Entrant)**.** When $\mathrm{CS}_{\mathrm{unc}} < \mathrm{CS}_0$, no cap clears (A) and the merger does not occur; realized consumer surplus is CS_0.

Proof. For every cap, $\mathrm{CS}(\bar{Q}, E) \leq \mathrm{CS}_{\mathrm{unc}} < \mathrm{CS}_0$ by (M1), so (A) fails. Every proposal is rejected and the merger blocked, so merging leaves the parties with Π_0, the same as not merging, and by the tie-breaking convention (footnote 20) they do not merge. □

B.3 Proof of Proposition 3.1

Proof. Working backward, at each pricing history the adding-up constraint has a unique root (Sections A.1 and A.2), which pins every firm's markup, so prices are the unique Bertrand equilibrium and the payoffs of Table 1 are the continuation values at the entry stage. By Assumption 3.1 the merger fails the initial screen, so it can proceed only with a remedy. If $\mathrm{CS}_{\mathrm{unc}} < \mathrm{CS}_0$ no capacity clears (A) (Proposition B.2), the approvable set is empty, and the merger does not occur with realized consumer surplus CS_0. Suppose $\mathrm{CS}_{\mathrm{unc}} \geq \mathrm{CS}_0$, so $Q_A \leq Q_{\mathrm{unc}}$ exists and by (M1) the approvable set is $[Q_A, \infty)$.

Realized profit is $\Pi^M_{\mathrm{m,ne}}$ at any capacity where entry fails and $\Pi^M(\bar{Q}, E)$ where it occurs, and $\Pi^M_{\mathrm{m,ne}} > \Pi^M(\bar{Q}, E)$ for every $\bar{Q} > 0$ at which entry occurs, since $\Pi^M(0, E) = \Pi^M_{\mathrm{m,ne}}$ (Section A.2) and Π^M is decreasing in $\bar{Q}$ (M3). The parties therefore prefer any approvable capacity that deters entry to every capacity that induces it, and (M3) selects among capacities only when all of them induce entry. By Proposition A.2(iv), (F) holds exactly on $[Q_F, Q_F^+]$, so every approvable capacity induces entry if and only if $[Q_A, \infty) \subseteq [Q_F, Q_F^+]$, that is if and only if $Q_F \leq Q_A$ and $Q_F^+ = \infty$, the second of which is $F_T \leq \pi^F_{\mathrm{unc}}$. This is the definition of entry being unavoidable, and the four cases below are exhaustive.

(i) Entry unavoidable and $Q_A \leq Q_{MP}$. Realized profit is $\Pi^M(\bar{Q}, E)$ at every approvable capacity, so by (M3) the parties propose the lowest, $\bar{Q}^* = Q_A$. Since $Q_A \leq Q_{MP}$ and Π^M is decreasing, $\Pi^M(Q_A, E) \geq \Pi^M(Q_{MP}, E) = \Pi_0$, so (MP) holds and the merger occurs. The entrant enters and $\mathrm{CS} = \mathrm{CS}(Q_A, E) = \mathrm{CS}_0$ by definition of Q_A.

(ii) Entry unavoidable and $Q_A > Q_{MP}$. Realized profit is again $\Pi^M(\bar{Q}, E)$ at every approvable

capacity, and by (M3) it is highest at Q_A; but $\Pi^M(Q_A, E) < \Pi^M(Q_{MP}, E) = \Pi_0$, so (MP) fails at every approvable capacity. At the proposal stage, therefore, any capacity below Q_A is strictly better: it is rejected, the merger is blocked, and the parties keep Π_0. All such capacities are equivalent, so the continuation value of merging is Π_0, the same as not merging. By the tie-breaking convention (footnote 20) the parties do not merge, and realized consumer surplus is CS_0.

(iii)–(iv) Entry avoidable. Some approvable capacity leaves $\pi^F < F_T$; at any such capacity realized profit is $\Pi^M_{\mathrm{m,ne}}$, which exceeds $\Pi^M(\bar{Q}, E)$ at every entry-inducing capacity and clears (MP) (Nocke and Schutz, 2025, Proposition 2). The parties propose one of them, the merger occurs, the entrant stays out, and realized consumer surplus is $\mathrm{CS}_{\mathrm{m,ne}} < \mathrm{CS}_0$ (Assumption 3.1). All such capacities deliver the same profit, so the proposal is not pinned down. The case is (iii), strategic deterrence, when $\mathcal{S} \neq \emptyset$, a capacity that would have induced entry while keeping the merger viable being available and not proposed. It is (iv), structural failure, when $\mathcal{S} = \emptyset$, no such capacity existing. The latter covers four configurations: no finite Q_F exists ($F_T > \pi^F_{\max}$), the corridor is empty ($Q_A > Q_{MP}$), or the two sets are disjoint, with $Q_F > Q_{MP}$ or $Q^+_F < Q_A$. □

B.4 Uncapped Access

Regime (c) of Proposition 3.2 uses the following.

Proposition B.3 (No SPE with Uncapped Entry)**.** Suppose $E \in \mathcal{A}$. Then no SPE exists in which the merger occurs and the entrant enters without a cap.

Proof. By definition of $\mathcal{A}$, $E < \bar{E}$. By Lemma A.1 and strict monotonicity of Π^M in E, $\Pi^M(Q_{\mathrm{unc}}, E) < \Pi_0$: uncapped entry renders the merger unprofitable, so no SPE has the merger together with entry at unbounded scale. CMS Proposition 5 is the case $E = 0$. □

B.5 Proof of Proposition 3.2

Proof. (a) Entry is not in Toyota's choice set without the remedy, so the relevant comparison at the unconditional screen is $\mathrm{CS}_{\mathrm{m,ne}}$ against CS_0; by Assumption 3.1 this fails, so the merger does not occur.

(b) Immediate from Proposition 3.1. When entry is unavoidable realized surplus is CS_0: at $\bar{Q}^* = Q_A$, where $\mathrm{CS}(Q_A, E) = \mathrm{CS}_0$ by definition of Q_A, if the corridor is nonempty, and at

the no-merger benchmark if it is not. When entry is avoidable the parties deter and realized surplus is $\mathrm{CS}_{\mathrm{m,ne}} < \mathrm{CS}_0$ by Assumption 3.1.

(c) The authority applies (A) to uncapped access, at which surplus is $\mathrm{CS}_{\mathrm{unc}}$. Against a weak entrant $\mathrm{CS}_{\mathrm{unc}} < \mathrm{CS}_0$, so the remedy is rejected and the merger does not occur (Proposition B.2). Let the entrant be strong. Removing the cap does not remove the entrant's participation constraint. If $F_T \leq \pi^F_{\mathrm{unc}}$ the entrant enters; by Assumption 3.2, $E < \bar{E}$, so uncapped entry leaves merged profit below the pre-merger level and (MP) fails (Proposition B.3), and the merger does not occur. If $F_T > \pi^F_{\mathrm{unc}}$ the entrant does not enter at any scale the regime permits; the merged entity then earns $\Pi^M_{\mathrm{m,ne}}$, which clears (MP) (Nocke and Schutz, 2025, Proposition 2), so the merger occurs and realized surplus is $\mathrm{CS}_{\mathrm{m,ne}}$ (Assumption 3.1). Proposition B.3 rules out an equilibrium with the merger and uncapped entry; it does not rule out the merger without entry.

(d) Contingent approval removes the no-entry continuation from the parties' choice set: a proposal that fails to draw entry leaves them at Π_0. They therefore propose a capacity in $\mathcal{S}$ whenever one exists and, by (M3), the smallest, $\max(Q_A, Q_F)$. Since CS is increasing (M1) and $\mathrm{CS}(Q_A, E) = \mathrm{CS}_0$, realized surplus is at least CS_0, strictly above it exactly when $Q_F > Q_A$. If $\mathcal{S} = \emptyset$ every approvable entry-inducing capacity violates (MP) and every other proposal leaves the parties at Π_0, so merging yields at most Π_0. By the tie-breaking convention (footnote 20) the parties do not merge, and realized surplus is CS_0.

(e) The authority maximizes CS over $\mathcal{S}$. By (M1) CS is strictly increasing in $\bar{Q}$, so the maximizer is $\max \mathcal{S}$, weakly above $\min \mathcal{S}$ and strictly above it when $\mathcal{S}$ is not a singleton. At $\max \mathcal{S}$ the capacity is approved and (MP) holds, with equality when $\max \mathcal{S} = Q_{MP}$, so the parties merge. If $\mathcal{S} = \emptyset$ no capacity satisfies (A), (MP) and (F), so the authority sets no cap and the merger would be blocked; by the tie-breaking convention (footnote 20) the parties do not merge. □

B.6 Ranking of the Regimes

Proposition B.4 (Ranking of Regimes)**.** Ranked by realized consumer surplus, (e) $\succeq$ (d) $\succeq$ (a) $\succeq$ (c) $\succeq$ (b). Each link is strict exactly as follows. The authority-set cap (e) is strictly above verified entry (d) when $\mathcal{S}$ contains more than one capacity. Verified entry is strictly above no remedy (a) when $\mathcal{S} \neq \emptyset$ and $Q_F > Q_A$. No remedy is strictly above uncapped access (c) when the entrant is strong and $F_T > \pi^F_{\mathrm{unc}}$. Uncapped access is strictly above the party-proposed cap (b) when the entrant is strong, $Q_F > Q_A$ and $F_T \leq \pi^F_{\mathrm{unc}}$. No remedy is

therefore strictly above the party-proposed cap exactly when the entrant is strong and entry is avoidable.

Proof. By Proposition 3.2, regime (a) delivers CS_0. Regime (b) delivers $\mathrm{CS}_{\mathrm{m,ne}} < \mathrm{CS}_0$ when the entrant is strong and entry is avoidable, and CS_0 otherwise. Regime (c) delivers $\mathrm{CS}_{\mathrm{m,ne}}$ when the entrant is strong and $F_T > \pi^F_{\mathrm{unc}}$, and CS_0 otherwise. Regime (d) delivers $\mathrm{CS}(\max(Q_A, Q_F), E) \geq \mathrm{CS}_0$ when $\mathcal{S} \neq \emptyset$ and CS_0 otherwise, and regime (e) delivers $\mathrm{CS}(\max \mathcal{S}, E)$ when $\mathcal{S} \neq \emptyset$ and CS_0 otherwise.

(e) and (d). Regime (e) delivers $\mathrm{CS}(\max \mathcal{S}, E)$, weakly above regime (d) by (M1) since $\max \mathcal{S} \geq \min \mathcal{S}$, and strictly above it when $\mathcal{S}$ is not a singleton, since $Q_{MP} < Q_{\mathrm{unc}}$ by Assumption 3.2 and (M3), so $\mathcal{S} \subseteq \mathcal{C}$ lies where CS is strictly increasing.

(d) and (a). Regime (d) rises strictly above CS_0 exactly when $\mathcal{S} \neq \emptyset$ and $Q_F > Q_A$: then $\max(Q_A, Q_F) = Q_F > Q_A$, and since CS is strictly increasing (M1) with $\mathrm{CS}(Q_A, E) = \mathrm{CS}_0$, it follows that $\mathrm{CS}(Q_F, E) > \mathrm{CS}_0$.

(a) and (c). Regime (c) falls below CS_0 exactly when the entrant is strong and $F_T > \pi^F_{\mathrm{unc}}$.

(c) and (b). Against a weak entrant both deliver CS_0. Let the entrant be strong. If $F_T > \pi^F_{\mathrm{unc}}$, every approvable capacity at or above Q_{unc} pays the entrant $\pi^F_{\mathrm{unc}} < F_T$, so entry is avoidable in (b) and both regimes deliver $\mathrm{CS}_{\mathrm{m,ne}}$. If $F_T \leq \pi^F_{\mathrm{unc}}$, regime (c) delivers CS_0, and regime (b) delivers $\mathrm{CS}_{\mathrm{m,ne}}$ exactly when entry is avoidable, which with $Q^+_F = \infty$ means $Q_F > Q_A$.

(a) and (b). Regime (b) falls below CS_0 exactly when the entrant is strong and entry is avoidable, so no remedy is strictly above the party-proposed cap exactly then. □

The level of the standard. The regimes change what approval requires and who selects the capacity, not the floor itself. Raising the floor is a further change. Nothing in the proof of Proposition 3.1(i) invokes CS_0 except through the definition of Q_A, since the parties choose the lowest approvable capacity whatever the floor. A higher floor therefore raises Q_A, and realized surplus with it, until Q_A passes Q_{MP} and the merger is abandoned. Up to that point the parties concede exactly the floor and no more.

C Counterfactual Specification: Toyota's Assumed Primitives

The four-scenario welfare comparison in Section 7 requires specifying primitives for Toyota's Proace City, which is not observed in the Chilean small commercial van segment during the sample period. We adopt the maintained assumptions described below, motivated by the institutional record (Section 2) and the fact that the Proace City is co-produced with the PSA platform models at the same manufacturing facility.

Utility. We use Toyota's observed Proace City product characteristics to estimate the valuation that consumers would have for the product. For ξ_F we use the equivalent average values of the models Berlingo and Partner. This assumption is justified by the institutional record, since the European Commission's investigation found that the Proace City is virtually identical to the Citroën Berlingo and Peugeot Partner in design and specification (European Commission, 2021).

Marginal cost. Toyota's marginal cost is estimated based on the supply estimation results. As with the utility specification, we use observed characteristics together with the average freight cost of the Berlingo and Partner models and their average $\hat{\omega}$, which we assign to ω_F. This reflects the co-production arrangement. Under the remedy, Toyota would import vehicles manufactured at the same PSA facility that produces the existing platform models. The remedy's cost-side provisions are modeled implicitly, since the fixed cost F_T bounded below in Section 7 is Toyota's effective fixed entry cost net of them, capturing logistics setup and distribution costs, dealer network setup, marketing, homologation, and other Chile-specific entry costs that those provisions were designed to offset. Unlike the merging parties, we do not allow for freight-related efficiencies.

Conduct. Under uncapped entry, Toyota plays Bertrand–Nash against the merged entity and against non-merging incumbents, while the merged entity uses its estimated post-merger conduct parameter $\hat{\kappa}$. Under the binding cap, Toyota sets price to clear exactly $\bar{Q}$ units (per Section 3), and the merged entity's first-order conditions incorporate Toyota's binding-quantity pricing through the demand system.

Sensitivity to ξ_F and ω_F. Because Toyota is not observed in the segment, Appendix S4 draws these primitives from the Partner alone and from the Berlingo alone and reports how

the diagnosis moves (Table S10). Table 7 in the body reports every identified object at the baseline.

D Bootstrap Procedure for Counterfactual Inference

We construct confidence intervals for counterfactual welfare quantities and the identified fixed-cost bounds via a parametric bootstrap that propagates sampling uncertainty in the BLP demand parameters through the counterfactual equilibrium calculations.

Bootstrap draws. Demand parameter draws are obtained as

$$\theta_D^{(b)} \sim \mathcal{N}\big(\hat{\theta}_D,\, \widehat{V}_D\big), \qquad b = 1, \ldots, B,$$

where $\hat{\theta}_D$ denotes the GMM point estimates from the demand estimation and $\widehat{V}_D$ their asymptotic covariance matrix from the second-step procedure. We draw $B = 200$ candidate parameter vectors.

What is perturbed versus held fixed. Demand parameters $(\hat{\beta}, \hat{\alpha}, \hat{\Sigma})$ are perturbed across draws. Marginal costs are re-recovered from the pricing first-order conditions at each draw and the cost-side coefficients γ are re-estimated on them, so γ varies across draws, while the conduct parameter is held fixed at $\hat{\kappa}$. Pinning κ isolates the sensitivity of the welfare quantities to demand estimation and to the cost side it induces, while conduct uncertainty is treated separately in the sensitivity tables. Toyota's counterfactual primitives (Appendix C) and the observed product characteristics are held fixed across draws.

Per-draw procedure. For each draw $\theta_D^{(b)}$:

1. Recompute the BLP inversion to obtain mean utilities $\delta_t(\theta_D^{(b)})$ that rationalize observed market shares.

2. Recompute implied marginal costs by inverting the supply first-order conditions at $\hat{\kappa}$, using the draw-specific demand derivatives.

3. For each of the four counterfactual scenarios (no-merger, merger with no entry, merger with capped entry, and merger with uncapped entry), solve for equilibrium prices by

iterating on the ζ-markup contraction of Morrow and Skerlos (2011) as implemented in PyBLP (Conlon and Gortmaker, 2020), until the largest price change is below 10^{-12}.

4. Compute the quantities of interest: consumer surplus by scenario, producer surplus, prices, profits, and the fixed-cost bound endpoints $\pi^F(Q^*, E)$ and $\pi^F_{\max}$.

Confidence intervals. We report percentile bootstrap confidence intervals: the 95% CI for any quantity of interest is the interval between its 2.5th and 97.5th percentiles across the valid draws. Because the procedure perturbs the demand parameters and re-estimates the cost side at each draw while holding the conduct parameter $\hat{\kappa}$ and Toyota's counterfactual primitives fixed, these intervals do not integrate over uncertainty in conduct or in Toyota's primitives. We therefore interpret them as partial statistical intervals. Sensitivity to conduct and to Toyota's primitives is reported separately as model uncertainty in Tables 13 (conduct), S10 (Toyota's primitives) and S11 (time-trend and market-size specification).

Convergence and screening. Under the cap, Toyota's marginal cost is scaled up until its volume meets the cap, with the scale found by Brent's method and prices re-solved at every trial value. A draw is discarded if that search fails, if its mean own-price elasticity falls outside $[-15, -1]$, or if the implied post-merger Toyota share is numerically zero. A third restriction is on the parameter space: a draw is discarded if it implies a negative merger efficiency, $E < 0$, since Lemma A.1 defines the counterfactual objects only on the admissible efficiency range. Because efficiency enters marginal cost through the interaction of freight cost with post-merger ownership, this is the restriction that the sampled coefficient be non-positive. Of the 200 draws, 11 violate it and 189 are retained for the percentile intervals. The convergence and plausibility screens bind for no draw, and both corridor endpoints Q_A and Q_{MP} are defined in every retained draw. Because the admissibility restriction removes draws in which the merger raises the merging parties' costs, the intervals are conditional on the maintained parameter space.

E Estimation Details

Demand

For a candidate parameter vector $\theta^D = (\beta, \alpha, \sigma_o)$ we compute, for each market t, the vector of mean utilities $\delta_t(\theta^D)$ that rationalizes observed shares. The inversion uses the Berry et al.

(1995) contraction,

$$\delta_{jt}^{(m+1)} = \delta_{jt}^{(m)} + \ln(s_{jt}^{obs}) - \ln(s_{jt}^{model}(\theta^D, \delta^{(m)})), \tag{15}$$

which under mild regularity conditions converges to the unique $\delta_t(\theta^D)$ replicating observed shares. Matching these to (2) yields the demand shocks

$$\xi_{jt}(\theta^D) = \delta_{jt}(\theta^D) - X_{jt}\beta + \alpha p_{jt}. \tag{16}$$

Parameters are estimated by GMM under $E[\xi_{jt}(\theta^D) \mid Z_{jt}] = 0$, where Z_{jt} collects the demand-side instruments. In the sample these generate

$$g_T(\theta^D) = \frac{1}{T}\sum_{t=1}^{T}\sum_{j=1}^{J_t} Z_{jt}'\,\xi_{jt}(\theta^D), \tag{17}$$

and the objective $Q_T(\theta^D) = g_T(\theta^D)' W_T\, g_T(\theta^D)$. Estimation proceeds in two stages. The first minimizes Q_T with $(Z'Z)^{-1}$ as the weighting matrix, on the instruments described in Section 5. From the resulting estimates we construct approximated optimal instruments, replacing Z_{jt} with the Jacobian of ξ_{jt} with respect to θ^D evaluated at the expected value of the demand shocks, zero (Chamberlain, 1987; Reynaert and Verboven, 2014). The second stage re-estimates on these moment conditions (Conlon and Gortmaker, 2020), and the reported estimates $\widehat{\theta}^D$ come from it.

Supply

Taking the demand estimates $\widehat{\theta}^D$ as given, the supply parameters solve the sample analogue of $E[\omega_{jt}(\gamma, \kappa)\, W_{jt}] = 0$:

$$g_T^S(\gamma, \kappa) = \frac{1}{T}\sum_{t=1}^{T}\sum_{j=1}^{J_t} W_{jt}'\,\omega_{jt}(\gamma, \kappa), \tag{18}$$

with objective $Q_T^S(\gamma, \kappa) = g_T^S(\gamma, \kappa)'\widetilde{W}_T\, g_T^S(\gamma, \kappa)$ and $\widetilde{W}_T$ positive semi-definite. We implement two-step GMM; the final estimates $(\widehat{\gamma}, \widehat{\kappa})$ minimize Q_T^S. Table 8 reports the full set of cost-shifter coefficients behind Table 5.

Table 8: Supply estimates, full cost-shifter set

	Single κ
Constant	-7075.152^{***}
	(489.900)
Diesel	856.197^{***}
	(189.009)
Weight	1.471^{***}
	(0.219)
Fuel economy	52.332^{**}
	(22.546)
Origin	2380.566^{***}
	(85.425)
Wheelbase	2.229^{***}
	(0.220)
Engine Size (L)	796.938^{***}
	(212.764)
Norm 6	555.682^{***}
	(93.900)
Freight price	1.067^{***}
	(0.192)
Efficiency (freight)	-0.612^{***}
	(0.202)
κ	0.394^{**}
	(0.193)
Demand spec	RC (origin)
κ reparametrization	Logistic
Market size	1.5× sales
Time controls	LQ trend

Notes: The Time controls row reports the controls included in the marginal-cost equation. Cost-shifter coefficients from GMM supply estimation (2-step). For logistic reparametrization (κ reparametrization = Logistic), reported κ is $\sigma(t) = e^t/(1+e^t)$ evaluated at the estimated logit-space parameter; standard errors via delta method. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

F Data Construction and Market Description

F.1 Sample Construction and Data Cleaning

Our main dataset is constructed from administrative information provided by the TGR. This agency is responsible for collecting green tax payments and records each vehicle sale together with a set of variables, including transaction price, model, brand, and detailed vehicle characteristics.

Before being commercialized in Chile, all new light and medium vehicles must obtain certification from the 3CV. As part of this process, the 3CV performs laboratory tests to determine pollutant emission levels, fuel efficiency, and compliance with national regulations governing safety, size, and vehicle components. Upon approval, each model is assigned a unique CIT code. This identifier allows us to link individual sales records in the TGR database to detailed technical specifications for each vehicle model.

We restrict the sample to small commercial van models. The data are available from January 2015 onward. In April 2023, the Euro 6 emissions standard became mandatory in Chile, requiring all vehicles sold to comply with this regulation. The introduction of this standard led to the exit of several models that were not upgraded to Euro 6 versions and, in some cases, to the exit of firms from the segment. It also entailed changes in vehicle production costs. To avoid these transitional effects from affecting our results, we restrict the sample to the period ending in March 2023.

We exclude electric models, whose prices are substantially higher than those of conventional vehicles and whose market share remains below 5% for most of the sample period. In addition, we drop models with very low sales volumes (less than 0.1% of total sales over the entire sample) and remove observations with implausibly low or high prices (below CLP 600,000 or above CLP 45 million). As noted above, each product is identified by a unique CIT code. However, we consolidate models that underwent rebadging processes.

F.2 Sales, Firms, and Models

Table 9: Sales, firms, and models of small commercial vans (Chile, 2015–2023)

Year	Sales	Firms		Models	
		All	Non-merging	All	Non-merging
2015	9,498	13	11	31	15
2016	10,540	13	11	36	20
2017	12,298	12	10	35	20
2018	10,039	12	10	33	18
2019	8,669	11	9	34	15
2020	5,997	10	8	28	14
2021	12,717	9	7	23	11
2022	9,354	8	6	23	9
2023	4,086	7	5	21	6
Total	83,198	#14	#12	#52	#25

Notes: 2023 covers January to March. # indicates the number of distinct firms or models over the whole period, which is not the sum of the yearly counts. Non-merging counts exclude FCA and PSA.

Table 10: Summary statistics

	All years		Pre-merger	Post-merger
Variable	Mean	SD	Mean	Mean
Price (thousands of CLP)	8,683.51	2,273.81	8,089.29	9,979.31
Weight (kg)	2,022.12	239.87	2,013.64	2,040.62
Wheelbase (mm)	2,718.17	142.12	2,698.27	2,761.57
Width (mm)	1,795.57	103.62	1,802.85	1,779.71
Engine size (L)	1.44	0.17	1.43	1.46
Performance (km/liter)	14.52	3.07	14.65	14.22
CO2 (gr/km)	151.99	20.09	150.92	154.33
Diesel (0–1)	0.63	0.48	0.63	0.63
EURO 6 (0–1)	0.11	0.31	0.04	0.25
Origin (0–1)	0.70	0.46	0.72	0.68
Observations	83,198		57,041	26,157

Notes: The table reports means and standard deviations of the main variables, weighted by units sold, for all years and for the pre-merger (2015–2020) and post-merger (2021–2023) periods separately.

Table 11: Market shares by firm – small commercial van sales (2019)

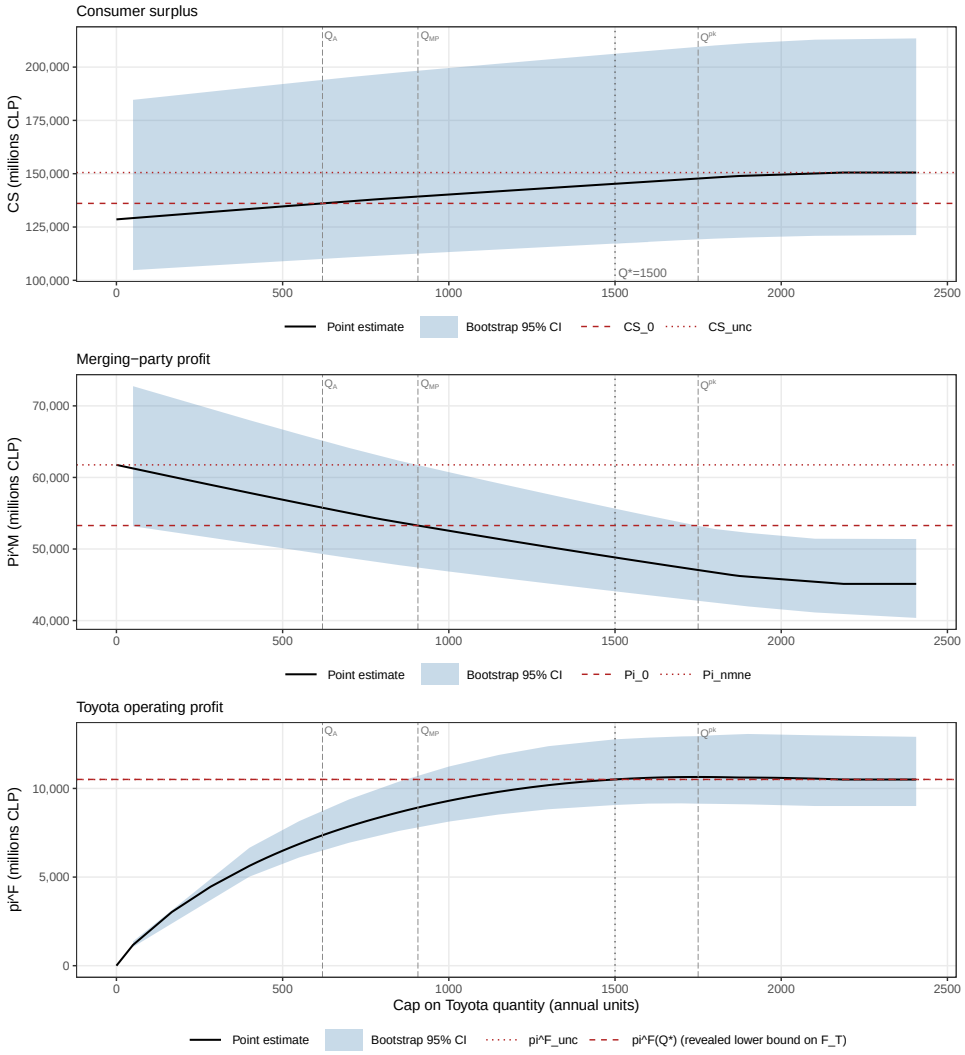
Firm	Market Share
PSA	0.426
FCA	0.179
Chevrolet	0.175
Renault-Nissan	0.084
Changan	0.038
Dongfeng	0.035
Mercedes-Benz	0.022
FAW	0.020
BAIC group	0.010
Suzuki	0.010
Lifan	0.001
HHI	2552.4

Notes: Market shares computed from quantity sold in 2019 (pre-merger baseline). HHI is the Herfindahl-Hirschman Index scaled to 10,000.

G Counterfactual Results

This appendix collects the counterfactual output Section 7 draws its numbers from: the welfare frontier across caps, the three conduct models and the diagnosis across them, the four market configurations, and the welfare decomposition.

Figure 4: Welfare Frontier Across Cap Levels



Notes: consumer surplus, merging-party profit and the entrant's operating profit against the annual cap $\bar{Q}$. Solid lines are point estimates and shaded bands 95% percentile bands across the 189 retained bootstrap draws. Long-dashed verticals mark Q_A, Q_{MP} and the cap Q^{pk} at which π^F peaks; the dotted vertical is the proposed 1,500-unit cap. Horizontal lines mark CS_0, Π_0, π^F_{unc}, and the revealed lower bounds on F_T. The corridor $[Q_A, Q_{MP}]$ lies entirely to the left of the region where π^F clears its lower bound, which is the disjointness in the text.

Table 12: Merger Effects under Different κ Values

	$\kappa = 0$	κ_{actual}	$\kappa = 1$	(1) − (2)	(3) − (2)
Prices (Merging, RAM)	10779.381	11673.831	13367.893	−894.450 (−7.7%)	1694.062 (14.5%)
Prices (Merging, Others)	11141.754	11820.304	12980.515	−678.550 (−5.7%)	1160.212 (9.8%)
Prices (Non Merging)	8389.963	8398.457	8415.014	−8.494 (−0.1%)	16.557 (0.2%)
Profit (Merging, RAM)	17801.157	19856.857	20757.169	−2055.700 (−10.4%)	900.312 (4.5%)
Profit (Merging, Others)	37669.318	41900.002	45497.174	−4230.684 (−10.1%)	3597.172 (8.6%)
Profit (Non Merging)	25227.708	27002.218	30077.690	−1774.510 (−6.6%)	3075.472 (11.4%)
Profit (Toyota, uncapped entry)	9544.527	10503.359	12065.113	−958.832 (−9.1%)	1561.754 (14.9%)
CS	140050.187	128579.586	112104.321	11470.601 (8.9%)	−16475.265 (−12.8%)
PS	80698.183	88759.077	96332.033	−8060.894 (−9.1%)	7572.957 (8.5%)
TS	220748.371	217338.663	208436.355	3409.708 (1.6%)	−8902.309 (−4.1%)

Notes: All three scenarios include merger efficiencies and, except for the Toyota row, describe the merger without entry; the Toyota row is its operating profit under uncapped entry at each conduct level. $\kappa = 0$: merged firm prices at Nash–Bertrand. κ_{actual}: estimated conduct parameter (observed). $\kappa = 1$: full internalization of cross-price effects. Differences (columns 4–5) are relative to κ_{actual}.

Table 13: The diagnosis across degrees of internalization

	$\kappa = 0$	$\hat{\kappa}$	$\kappa = 1$
$\text{CS}_{m,ne}/\text{CS}_0 - 1$	+2.9%	−5.5%	−17.6%
Assumption 3.1 (harm)	fails	holds	holds
Q_A (lowest approvable cap)	all caps	620	1,694
Q_{MP} (highest merger-viable cap)	–	907	1,422
Corridor $\mathcal{C}$	–	[620, 907]	$\emptyset$ ($Q_A > Q_{MP}$)
$\pi^F(Q^*, E)$ at the proposed cap	–	10,509.0	12,146.4
Peak of π^F	–	10,646.7 at 1,750	12,414.0 at 1,825
$\mathcal{S} = \emptyset$?	–	yes	yes
Outcome	merger clears the screen	structural failure	merger fails outright

Notes: point estimates. The no-entry surplus is the Table 12 source; the corridor objects come from the cap frontier solved at each conduct. At $\kappa = 0$ the merger without entry leaves consumers above CS_0, so every cap is approvable and no remedy is required; no frontier is needed. At $\kappa = 1$ the lowest approvable cap exceeds the highest merger-viable one, so no cap clears both constraints and the merger fails outright, before the entrant's participation constraint enters. Profits in millions of CLP over January 2021–March 2023; capacities in annual units.

Table 14: Welfare across the four market configurations

	1 No Merge	2 Merge (No Entry)	3 Merge (Entry No Cap)	4 Merge (Entry Cap)	2-1	3-1	4-1
Selected by regime	*(a), (d) and (e)*	*(b) and (c), no entry*	*off path*	*off path*			
Prices (Merging, RAM)	11070.171	11673.831	11096.674	11211.419	603.661 (5.5%)	26.503 (0.2%)	141.248 (1.3%)
Prices (Merging, Others)	11361.198	11820.304	11075.902	11220.310	459.106 (4.0%)	−285.296 (−2.5%)	−140.888 (−1.2%)
Prices (Non Merging)	8393.471	8398.457	8383.022	8386.036	4.986 (0.1%)	−10.449 (−0.1%)	−7.435 (−0.1%)
Profit (Merging, RAM)	16586.248	19856.857	15639.662	16504.433	3270.608 (19.7%)	−946.586 (−5.7%)	−81.816 (−0.5%)
Profit (Merging, Others)	36694.784	41900.002	29488.855	32318.427	5205.218 (14.2%)	−7205.929 (−19.6%)	−4376.357 (−11.9%)
Profit (Non Merging)	26015.819	27002.218	24023.642	24632.579	986.399 (3.8%)	−1992.177 (−7.7%)	−1383.240 (−5.3%)
Prices (Toyota)	—	—	10967.523	11628.531	—	—	—
Profit (Toyota)	—	—	10503.359	10509.026	—	—	—
CS	136077.569	128579.586	150542.726	145268.118	−7497.982 (−5.5%)	14465.158 (10.6%)	9190.550 (6.8%)
PS	79296.852	88759.077	79655.519	83964.465	9462.225 (11.9%)	358.667 (0.5%)	4667.613 (5.9%)
TS	215374.420	217338.663	230198.245	229232.583	1964.243 (0.9%)	14823.825 (6.9%)	13858.163 (6.4%)

Notes: Prices are in thousands of CLP. Profits and welfare measures are in millions of CLP, accumulated over the post-merger period (January 2021–March 2023). Scenario 1 (No Merge): $\kappa = 0$ throughout, marginal costs without merger efficiencies. Scenario 2 (Merge, No Entry): the observed post-merger equilibrium at $\hat{\kappa} = 0.394$, Toyota absent. Scenario 3 (Merge, Entry): uncapped Toyota entry, Bertrand–Nash against the merged entity at $\hat{\kappa}$. Scenario 4 (Merge, Entry Cap): entry under the 1,500-unit annual ceiling, which binds, so Toyota prices to clear the cap. Columns 2–1, 3–1, and 4–1 report differences relative to Scenario 1, with percentage changes in parentheses. The second row records which regime of Proposition 3.2 selects each configuration; Scenarios 3 and 4 are selected by none at the estimates; Scenario 4 is the configuration the authority credited at approval. CS: consumer surplus. PS: producer surplus. TS: total surplus.

Table 15: Welfare decomposition (unilateral effects vs. efficiencies)

	Merge	No Merge	Unilateral	Efficiencies	Total
Prices (Merging, RAM)	11673.831	11070.171	869.352 (7.9%)	−265.692 (−2.4%)	603.661 (5.5%)
Prices (Merging, Others)	11820.304	11361.198	646.688 (5.7%)	−187.582 (−1.7%)	459.106 (4.0%)
Prices (Non Merging)	8398.457	8393.471	8.497 (0.1%)	−3.511 (−0.0%)	4.986 (0.1%)
Profit (Merging, RAM)	19856.857	16586.248	1815.992 (10.9%)	1454.617 (8.8%)	3270.608 (19.7%)
Profit (Merging, Others)	41900.002	36694.784	3926.025 (10.7%)	1279.193 (3.5%)	5205.218 (14.2%)
Profit (Non Merging)	27002.218	26015.819	1716.230 (6.6%)	−729.831 (−2.8%)	986.399 (3.8%)
CS	128579.586	136077.569	−10660.517 (−7.8%)	3162.534 (2.5%)	−7497.982 (−5.5%)
PS	88759.077	79296.852	7458.247 (9.4%)	2003.978 (2.5%)	9462.225 (11.9%)
TS	217338.663	215374.420	−3202.270 (−1.5%)	5166.513 (2.4%)	1964.243 (0.9%)

Notes: Prices are in thousands of CLP. Profits and welfare measures are in millions of CLP, accumulated over the post-merger period (January 2021–March 2023). The "No Merge" column reports the counterfactual equilibrium with $\kappa = 0$ throughout and marginal costs at their estimated post-merger values net of the merger-specific efficiency. The "Unilateral" column isolates the price effect of the conduct shift, the difference between the equilibria at $\hat{\kappa}$ and at $\kappa = 0$, both at those same marginal costs. The "Efficiencies" column captures the partial offsetting effect of merger-specific cost reductions, recovered from the estimated efficiency parameter in the supply specification. The "Total" column is the sum of the two effects and equals the difference between the "Merge" and "No Merge" columns. Percentage changes in parentheses are relative to the "No Merge" baseline. CS: consumer surplus. PS: producer surplus. TS: total surplus.

Supplementary Appendix

The material below supports the results in the paper but is not required to follow the argument. It collects robustness for the reduced-form design, alternative demand and conduct specifications, and further counterfactual output. The derivations and proofs the propositions rest on, the imputation of the entrant's primitives, the bootstrap procedure, the estimation details, the construction of the estimation sample, and the counterfactual tables the text draws on are in the Appendix.

Contents

S1	Reduced-Form Robustness	65
S2	Demand Estimation Robustness	67
S3	Conduct Estimation Robustness	69
S4	Additional Counterfactual Results	75

S1 Reduced-Form Robustness

Table S1: DiD estimates: FCA-PSA vs. non-merging firms, by treatment date (October 2019 announcement, January 2021 consummation)

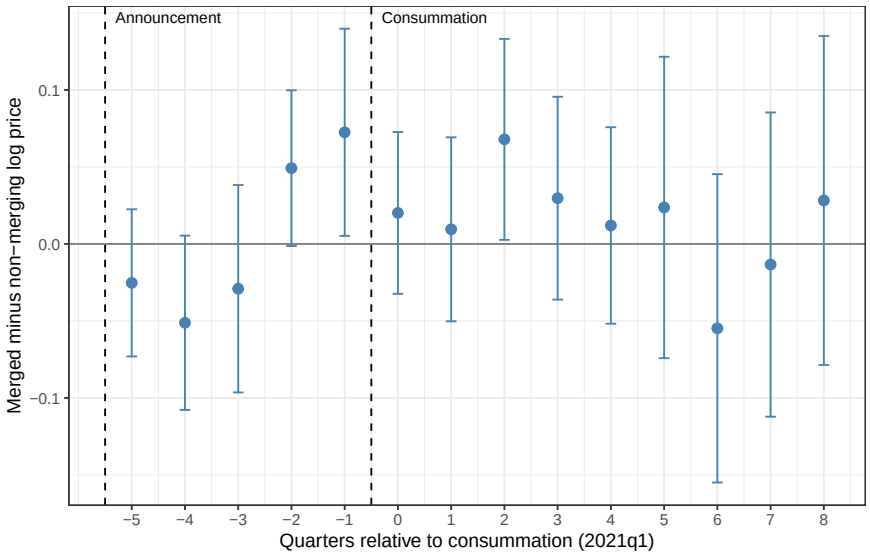
	Post $\geq$ 2019m10	Post $\geq$ 2021m1
Panel A: Overall		
β_{DiD} (Merged $\times$ Post)	0.007	0.015
	(0.025)	(0.027)
Panel B: By product group		
Merging (RAM)	0.005	0.021
	(0.027)	(0.030)
Merging (Others)	0.012	0.008
	(0.039)	(0.037)
Observations	2248	2248

Notes: Two-way FE (product and time). Clustered SE by product in parentheses. Panel A reports the baseline DiD coefficient for all PSA-FCA products jointly. Panel B reports group-specific coefficients from a separate regression that disaggregates the merged portfolio: Merging (RAM) pools RAM V700 City and RAM V1000; Merging (Others) pools RAM V700 Rapid and PSA platform models (Partner, Berlingo, Combo). Baseline category in Panel B: non-merging firms. * p<0.10, ** p<0.05, *** p<0.01.

S1.1 Product Entry and Exit

Group averages of log prices move when products enter or exit, and when a product sells in only some months of a quarter, so the faster rise of the non-RAM merged products in Figure 2 need not reflect prices. Figure S1 removes both effects. It estimates the difference-in-differences of Section 4.2 quarter by quarter, with product and month fixed effects, against a reference period that pools every quarter before the announcement. Merged prices do not step up relative to non-merging prices after consummation. The coefficients fluctuate around zero, with a temporary rise in the two quarters before consummation that reverses in 2021q1. Because 2020q4 falls within that rise, figures normalized to 2020q4 place merged prices below non-merging prices in the surrounding quarters.

Figure S1: Merged versus non-merging prices by quarter



Notes: Coefficients on Merged × quarter relative to consummation (2021q1) from a regression of log price on product and month fixed effects. The reference period pools all quarters before the announcement quarter (2019q4). Bars are 95% confidence intervals from product-clustered standard errors.

S2 Demand Estimation Robustness

Table S2: Demand estimates, full coefficient set

	Logit	RC (log-normal)	NL (one nest)	NL (origin)	RC (origin)
Price[a]	−0.844***	−0.249	−0.024***	−0.274***	−0.533***
	(0.059)	(1.399)	(0.006)	(0.060)	(0.038)
Wheelbase	0.001***	0.001**	0.000**	0.000**	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Weight	0.001***	0.001***	0.000***	0.000***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Performance (km/liter)	0.267***	0.270***	0.007***	0.095***	0.196***
	(0.036)	(0.040)	(0.002)	(0.024)	(0.034)
Engine Size (L)	0.725***	0.706	0.017*	0.325***	0.724***
	(0.252)	(0.645)	(0.010)	(0.124)	(0.253)
Diesel	0.954***	0.989*	0.027***	0.360***	0.735***
	(0.203)	(0.526)	(0.009)	(0.108)	(0.195)
Origin	1.975***	2.083*	0.057***	1.089***	−4.583***
	(0.189)	(1.148)	(0.015)	(0.105)	(1.610)
CO2	0.010**	0.010**	0.000*	0.003	0.006
	(0.004)	(0.005)	(0.000)	(0.002)	(0.004)
Norm 6	1.136***	1.152***	0.031***	0.463***	0.984***
	(0.150)	(0.343)	(0.009)	(0.106)	(0.145)
ρ (nesting)			0.973***	0.571***	
			(0.007)	(0.082)	
σ_{constant}		88.327***			
		(0.626)			
σ_{price}		0.231			
		(3.415)			
σ_{origin}					10.820***
					(2.917)
Mean Own Price Elast	−7.139	−7.276	−7.483	−5.236	−4.416
Market size	1.5× sales	1.5× sales	1.5× sales	1.5× sales	1.5× sales

Notes: Price is measured in millions of CLP. All estimations include year and month-of-year FE. [a]For log-normal RC columns, the reported price coefficient is μ (log-mean); the implied median is $-\exp(\mu)$.

Table S3 varies the market-size assumption. The potential market is defined as recorded sales scaled by $(1 + r)$, and the columns report $r = 0.5$, $r = 1$ and $r = 2$, that is a potential market 50%, 100% and 200% larger than observed sales. The assumption shifts the outside-good share and hence the level of mean utilities recovered from the inversion, but the price coefficient and the implied elasticities remain stable across the three. Market size enters only through the outside share, so $\hat{\alpha}$ is unchanged across the three while $\hat{\sigma}_o$ and the mean origin coefficient grow together, and the share of consumers who prefer manufacture outside China moves from 34% to 26%. The market-size assumption rescales the origin dimension without changing it.

Table S3: Demand estimates under alternative market sizes

	RC (origin)	RC (origin)	RC (origin)
Price[a]	−0.533***	−0.533***	−0.533***
	(0.038)	(0.038)	(0.038)
Wheelbase	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)
Weight	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)
Performance (km/liter)	0.196***	0.196***	0.196***
	(0.034)	(0.034)	(0.034)
Engine Size (L)	0.724***	0.724***	0.724***
	(0.253)	(0.252)	(0.252)
Diesel	0.735***	0.734***	0.734***
	(0.195)	(0.195)	(0.195)
Origin	−4.583***	−7.098***	−12.470***
	(1.610)	(2.408)	(4.104)
CO2	0.006	0.006	0.006
	(0.004)	(0.004)	(0.004)
Norm 6	0.984***	0.984***	0.984***
	(0.145)	(0.144)	(0.144)
σ_{origin}	10.820***	13.491***	19.187***
	(2.917)	(3.639)	(5.305)
Mean Own Price Elast	−4.416	−4.435	−4.450
Market size	1.5× sales	2× sales	3× sales

Notes: Price is measured in millions of CLP. All estimations include year and month-of-year FE.

S3 Conduct Estimation Robustness

Table S4 varies the time controls in the marginal-cost equation. Column (1) is the baseline, with linear and quadratic trends. Column (2) drops the time controls, and column (3) uses year fixed effects instead. The baseline gives the smallest estimate, $\hat{\kappa} = 0.394$, against 0.759 without time controls and 0.440 with year fixed effects. Flexible time controls cost precision,

because they absorb time-series variation that helps identify conduct.

Table S4: Supply Estimation Under Different Time Controls

	Single κ	Single κ	Single κ
Constant	−7075.152***	−8554.470***	−7991.811***
	(489.900)	(537.265)	(481.322)
Diesel	856.197***	742.688***	497.747***
	(189.009)	(209.487)	(184.766)
Weight	1.471***	1.181***	1.770***
	(0.219)	(0.240)	(0.216)
Fuel economy	52.332**	100.458***	96.519***
	(22.546)	(25.556)	(21.540)
Origin	2380.566***	2322.640***	2573.136***
	(85.425)	(93.796)	(85.458)
Wheelbase	2.229***	2.315***	2.138***
	(0.220)	(0.222)	(0.213)
Engine Size (L)	796.938***	1083.117***	459.729**
	(212.764)	(229.792)	(211.745)
Norm 6	555.682***	690.106***	461.030***
	(93.900)	(108.072)	(83.114)
Freight price	1.067***	2.667***	1.599***
	(0.192)	(0.212)	(0.204)
Efficiency (freight)	−0.612***	−1.758***	−1.197***
	(0.202)	(0.212)	(0.198)
κ	0.394**	0.759***	0.440**
	(0.193)	(0.087)	(0.208)
Demand spec	RC (origin)	RC (origin)	RC (origin)
κ reparametrization	Logistic	Logistic	Logistic
Market size	1.5× sales	1.5× sales	1.5× sales
Time controls	LQ trend	None	Year FE

Notes: The Time controls row reports the controls included in the marginal-cost equation. Cost-shifter coefficients from GMM supply estimation (2-step). For logistic reparametrization (κ reparametrization = Logistic), reported κ is $\sigma(t) = e^t/(1+e^t)$ evaluated at the estimated logit-space parameter; standard errors via delta method. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table S5 varies the market size, with the potential market equal to recorded sales times $(1 + r)$ for $r \in \{0.5, 1, 2\}$ and $r = 0.5$ the baseline. The conduct estimate rises with market size and stays significant at the 5% level.

Table S5: Supply Estimation Under Different Market Sizes

	Single κ	Single κ	Single κ
Constant	−7075.152***	−7087.086***	−7086.319***
	(489.900)	(484.643)	(481.178)
Diesel	856.197***	866.316***	875.291***
	(189.009)	(186.939)	(185.572)
Weight	1.471***	1.505***	1.528***
	(0.219)	(0.217)	(0.216)
Fuel economy	52.332**	54.184**	55.373**
	(22.546)	(22.302)	(22.131)
Origin	2380.566***	2321.680***	2280.042***
	(85.425)	(83.937)	(83.071)
Wheelbase	2.229***	2.227***	2.222***
	(0.220)	(0.219)	(0.218)
Engine Size (L)	796.938***	819.069***	836.104***
	(212.764)	(210.244)	(208.755)
Norm 6	555.682***	629.007***	677.897***
	(93.900)	(92.022)	(91.393)
Freight price	1.067***	0.973***	0.907***
	(0.192)	(0.188)	(0.186)
Efficiency (freight)	−0.612***	−0.642***	−0.664***
	(0.202)	(0.201)	(0.200)
κ	0.394**	0.417**	0.436**
	(0.193)	(0.196)	(0.198)
Demand spec	RC (origin)	RC (origin)	RC (origin)
κ reparametrization	Logistic	Logistic	Logistic
Market size	1.5× sales	2× sales	3× sales
Time controls	LQ trend	LQ trend	LQ trend

Notes: The Time controls row reports the controls included in the marginal-cost equation. Cost-shifter coefficients from GMM supply estimation (2-step). For logistic reparametrization (κ reparametrization = Logistic), reported κ is $\sigma(t) = e^t/(1+e^t)$ evaluated at the estimated logit-space parameter; standard errors via delta method. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table S6 re-estimates the supply side taking the nested logit on origin as the demand system, the specification behind the footnote in Section 7. The freight and efficiency coefficients are nearly unchanged, and the conduct parameter is higher but inside the baseline's 95% interval.

Table S6: Supply Estimation Under an Alternative Demand Specification

	Single κ	Single κ
Constant	−7075.152***	−5651.977***
	(489.900)	(471.190)
Diesel	856.197***	1268.592***
	(189.009)	(179.518)
Weight	1.471***	1.270***
	(0.219)	(0.212)
Fuel economy	52.332**	15.832
	(22.546)	(21.425)
Origin	2380.566***	2478.026***
	(85.425)	(86.533)
Wheelbase	2.229***	1.889***
	(0.220)	(0.219)
Engine Size (L)	796.938***	1173.575***
	(212.764)	(206.712)
Norm 6	555.682***	291.848***
	(93.900)	(104.813)
Freight price	1.067***	0.977***
	(0.192)	(0.191)
Efficiency (freight)	−0.612***	−0.599***
	(0.202)	(0.200)
κ	0.394**	0.564***
	(0.193)	(0.197)
Demand spec	RC (origin)	NL (origin)
κ reparametrization	Logistic	Logistic
Market size	1.5× sales	1.5× sales
Time controls	LQ trend	LQ trend

Notes: The Time controls row reports the controls included in the marginal-cost equation. Cost-shifter coefficients from GMM supply estimation (2-step). For logistic reparametrization (κ reparametrization = Logistic), reported κ is $\sigma(t) = e^t/(1+e^t)$ evaluated at the estimated logit-space parameter; standard errors via delta method. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

These exercises hold the specification of merger efficiencies fixed, and Table S7 varies it. Column (1) is the baseline, in which merger-specific cost changes enter through freight. Column (2) adds a post × merged level term, a post-merger cost shift that does not scale with freight, and columns (3) and (4) repeat both specifications with the exchange-rate, steel and aluminum input-price indices added to the freight price. With the level term, the freight slope flips sign and loses significance (from -0.612, $t = -3.0$, to 0.351, $t = 0.6$), $\hat{\kappa}$ rises from 0.394 to 0.899, and the implied efficiency rises from 3.4% to 33.9% of marginal cost. With the input-price indices added and no level term, column (3), the estimate is smaller and no longer significant at the 5% level.

From prices alone, a post-merger fall in marginal cost and a rise in internalization are observationally close, and when both are allowed the level term absorbs the post-merger cost change. What separates κ from a post-merger cost shift in the baseline is therefore the restriction that merger-specific cost changes operate through freight. The merging parties source the relevant models from a common production facility and import every unit sold,

so freight is the economically motivated channel.

Table S7: Conduct and Efficiency Under Alternative Cost Specifications

	(1) Baseline	(2) + level term	(3) + input-price indices	(4) + both
$\hat{\kappa}$	0.394**	0.899***	0.280	0.744***
	(0.193)	(0.163)	(0.213)	(0.196)
Efficiency: freight slope	−0.612***	0.351	−0.742***	0.047
	(0.202)	(0.548)	(0.226)	(0.495)
Efficiency: post × merged level	–	−2526.0***	–	−1943.0***
	(–)	(313.9)	(–)	(278.0)
Implied efficiency (%)	−3.4	−33.9	−4.0	−25.6
Input-price indices	Freight	Freight	Freight, FX, steel, aluminum	Freight, FX, steel, aluminum

Notes: Column (1) is the baseline marginal-cost specification, in which merger efficiencies enter as a freight-proportional shift (post × merged × freight price). Columns (2) and (4) add a post × merged *level* term, relaxing the restriction that the post-merger cost change scales with freight. Columns (3) and (4) add the exchange-rate, steel and aluminum input-price indices to the freight price. Standard errors in parentheses; $\hat{\kappa}$ is in conduct space, delta-method corrected. Implied efficiency is the mean log change in marginal cost for merged products post-merger. Adding the level term flips the sign of the freight slope, leaves it insignificant, and roughly doubles $\hat{\kappa}$; the freight restriction is what separates κ from a post-merger cost shift in the baseline. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table S8 reports weak-identification diagnostics for κ. The numerical derivative of the recovered marginal cost with respect to κ, evaluated at $\hat{\kappa}$, is treated as an endogenous regressor instrumented by the differentiation instruments. Across the baseline, the alternative time controls and the nested logit on origin, the effective first-stage F ranges from 67 to 85, against Montiel Olea–Pflueger 5% critical values of at most 15.8 for a worst-case bias of 10%.

Table S8: Weak-Identification Diagnostics for the Conduct Parameter

Specification	$\hat{\kappa}$ (SE)	First-stage F	KP robust F	Effective F	Critical value
Baseline (LQ trend)	0.394 (0.193)	114.9	136.1	83.0	15.7
No time trend	0.759 (0.087)	92.5	107.2	67.2	15.8
Year fixed effects	0.440 (0.208)	109.4	134.5	81.5	15.8
Nested logit (origin)	0.564 (0.197)	116.4	140.7	84.6	15.7

Notes: $\hat{\kappa}$ and its delta-method standard error are from the nonlinear GMM supply estimation. The first-stage F is the Cragg–Donald statistic, the KP robust F the Kleibergen–Paap heteroskedasticity-robust statistic, and the effective F the Montiel Olea and Pflueger (2013) statistic, our preferred measure with many instruments. The critical value is the Montiel Olea–Pflueger 5% value for a worst-case bias of 10%, computed for each row from its effective degrees of freedom; the Staiger–Stock rule of thumb for the first-stage F is 10. Excluded instruments are the own- and rival-differentiation instruments described in Section 5.

Conduct testing. The exercises above ask how $\hat{\kappa}$ moves and whether it is identified. The conduct test asks which values of κ the data can reject. Table S9 reports the moment-based

Rivers–Vuong test of Backus et al. (2021), which compares each fixed κ with post-merger Bertrand–Nash on a single scalar moment, the covariance between the cost residual that κ implies and the markup difference predicted from the instruments, $g(z) = \mathbb{E}[\Delta\eta \mid z]$. Conduct moves markups only on the merging parties' post-merger observations, 277 of 2,248 from 17 products, so the test rests on those.

The cost function is the linear specification of the estimation. The first stage regresses the markup difference on the cost regressors and on the differentiation instruments interacted with the merging parties' post-merger indicator, so the cost regressors enter g as they enter the cost function and only excluded variation identifies the test. In the first column the instruments enter linearly. In the second they enter as a cubic polynomial, following Calder-Wang and Kim (2024). The bootstrap repeats the whole procedure with clusters at the product-year level, the counterpart of the clustering in Backus et al. (2021).

The linear first stage separates no κ from Nash ($|T| = 0.56$ at $\kappa = 1$). The excluded instruments then explain 0.25 of the markup difference net of the cost regressors, against 0.76 with the cubic. The cubic first stage rejects full internalization ($|T| = 2.72$, and 2.36 when the first stage is cross-fitted by product-year) and does not separate $\hat{\kappa}$ from Nash ($|T| = 1.49$). At $\kappa = 0.1$ it favors partial internalization over Nash ($T = 2.02$). Neither set is a confidence interval for κ, which is [0.02, 0.77] (Section 6).

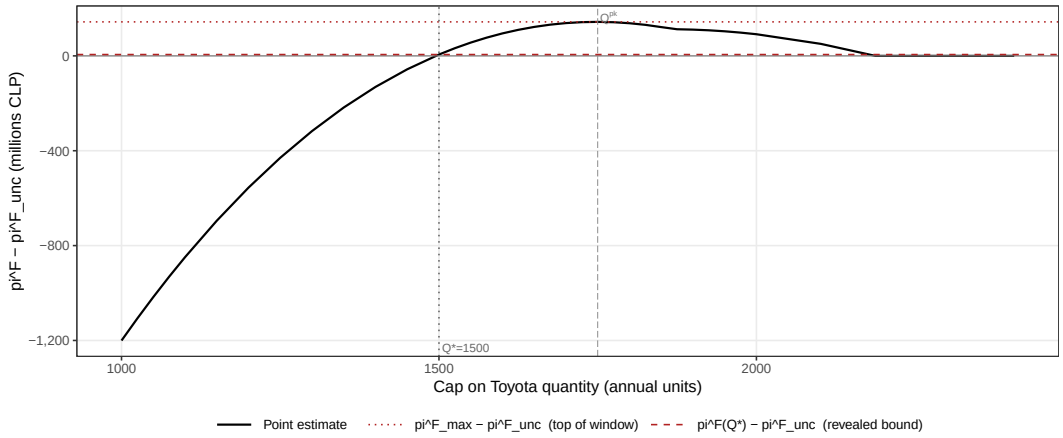
Table S9: Conduct testing: Rivers–Vuong separation from Bertrand–Nash

	First stage	
κ	Linear	Cubic
0.10	1.84	2.02*
0.20	1.81	1.91
0.30	1.77	1.75
0.40	1.73	1.47
0.50	1.66	0.96
0.60	1.56	0.09
0.70	1.37	−1.08
0.80	0.97	−2.03*
0.90	0.22	−2.51*
1.00	−0.56	−2.72*
$\hat{\kappa} = 0.394$	1.73	1.49
Not separated from $\kappa = 0$	[0, 1.0]	[0.2, 0.7]
Excluded terms	28	56
Partial R^2 at $\kappa = 1$	0.25	0.76

Notes: Rivers–Vuong statistics of each κ against post-merger Bertrand–Nash ($\kappa = 0$), following Backus et al. (2021). Negative values favor Bertrand–Nash, positive values favor the row's κ, and * marks $|T| > 1.96$. The partial R^2 is the share of the markup difference, net of the cost regressors, that the excluded terms explain. Standard errors come from 1,500 bootstrap replications clustered by product and year (264 clusters). Conduct moves markups on 277 of 2,248 observations.

S4 Additional Counterfactual Results

Figure S2: The entrant's operating profit relative to unconstrained entry



Notes: the object behind the revealed-preference bound. The curve is the entrant's operating profit π^F of the third panel of Figure 4, measured from its unconstrained level π^F_{unc} and shown for caps above 1,000 units, because the interior peak is about 1.4% of π^F and is invisible on the levels axis. π^F rises, crosses the unconstrained level below the proposed cap, peaks at Q^{pk}, declines, and flattens. The bound non-entry delivers at the proposed cap is therefore not the entrant's best available profit, and the vertical distance between the two horizontal lines is the unidentified window, which runs up to $\pi^F_{\max}$.

Robustness to Toyota's assumed primitives. Because Toyota is not in the sample, the classification rests on the demand and cost primitives we impute to the Proace City, which the baseline sets to the average of the Berlingo and Partner models (Appendix C). Table S10 reports the diagnosis when they are drawn from a single comparable model instead. The strong-entrant classification holds in every case, with Δ_{class} from +7.32% under the Berlingo-only primitives to +14.55% under the Partner-only primitives, against 10.63% at the baseline. The merger-without-entry harm ($\Delta_{\text{remedy}} = -5.51\%$) and the no-entry profit gain ($\Delta\Pi^M = 15.91\%$) do not involve Toyota and are identical across scenarios by construction. What moves is Toyota's volume relative to the cap, which the counterfactual imposes year by year. It exceeds 1,500 in every full year under the baseline, Partner-only and with-efficiencies primitives, and in one of the two full years under Berlingo-only, so the cap binds and non-entry identifies $F_T > \pi^F(Q^*, E)$ in each. Entrant profit at Q_{MP} falls short of that bound in every row, so $\mathcal{S} = \emptyset$ throughout.

Table S10: Sensitivity of the entrant diagnosis to Toyota's assumed primitives

Scenario	Δ_{class}	q^{unc}	Q_A	Q_{MP}	$\pi^F(Q_{MP})$	$\mathcal{S} = \emptyset$?	Diagnosis
Baseline (Partner–Berlingo average)	+10.63%	1,869	620	907	8,920	yes	strong, binds every year
Partner only	+14.55%	2,463	630	912	10,610	yes	strong, binds every year
Berlingo only	+7.32%	1,432	610	911	7,557	yes	strong, not every year
Baseline with efficiencies	+11.79%	2,015	622	897	9,388	yes	strong, binds every year

Notes: point estimates. Cap = 1,500 units. q^{unc} is Toyota's unconstrained annual volume, the minimum across full calendar years, since the counterfactual imposes the ceiling in each year; the cap binds in every year when it exceeds the cap. $\pi^F(Q_{MP})$ is the entrant's profit at the top of the corridor, interpolated on each scenario's frontier, and $\mathcal{S} = \emptyset$ reads yes when it falls short of $\pi^F(Q^*, E)$, the bound non-entry reveals. Under the Berlingo-only primitives the proposed cap lies past the peak of π^F, so the comparison there is not mechanical; it holds because Q_{MP} sits far below the cap.

Robustness to the estimation specification. Holding Toyota's primitives at the baseline imputation, Table S11 traces the diagnosis across 4 specifications: the time controls in the marginal-cost equation (the baseline linear–quadratic trend and year fixed effects) and the market size (the outside-share scaling, doubled and tripled). The strong-entrant classification holds in every row, with Δ_{class} between 10.63% and 12.06%; Toyota's unconstrained volume runs 1,788–1,869 units, above the cap in every row, so non-entry delivers $F_T > \pi^F(Q^*, E)$; and the merger-without-entry harm stays between −6.56% and −4.19%.

What moves is the corridor. Q_{MP} ranges from 907 at the baseline to 1,255 under year fixed effects, which also lower Q_A to 442 and raise the no-entry profit gain to +22.10%; doubling and tripling market size move Q_{MP} to 966 and 1,014. The corridor is therefore sensitive to the same time controls that move $\hat{\kappa}$ (Appendix S3). The verdict is not. No specification places the 1,500-unit cap inside $[Q_A, Q_{MP}]$, and entrant profit at Q_{MP}, from 8,920 at the baseline to 9,686 under year fixed effects, falls short of each specification's own revealed bound $\pi^F(Q^*, E)$, so $\mathcal{S} = \emptyset$ in every row, by a margin of 1,589 at the baseline and 335 under year fixed effects.

Table S11: Robustness of the entrant diagnosis to the estimation specification

Specification	Δ_{class}	q^{unc}	Q_A	Q_{MP}	$\pi^F(Q_{MP})$	Cap $\in [Q_A, Q_{MP}]$?	$\mathcal{S} = \emptyset$?	Δ_{remedy}	$\Delta\Pi^M$	Diagnosis
Time trend										
Linear–quadratic (baseline)	+10.63%	1,869	620	907	8,920	no	yes	−5.51%	+15.91%	strong, binds every year
Year fixed effects	+12.06%	1,788	442	1,255	9,686	no	yes	−4.19%	+22.10%	strong, binds every year
Market size										
×2 scaling	+11.05%	1,836	637	966	8,957	no	yes	−6.10%	+16.23%	strong, binds every year
×3 scaling	+11.31%	1,813	651	1,014	8,981	no	yes	−6.56%	+16.47%	strong, binds every year

Notes: point estimates. Cap = 1,500 units. Every row holds Toyota's primitives at the baseline imputation and varies only the time trend or the market size. $\pi^F(Q_{MP})$ is the entrant's profit at the top of the corridor, interpolated on the welfare frontier; since π^F rises across the corridor in every row, $\mathcal{S} = \emptyset$ whenever it falls short of the revealed bound $\pi^F(Q^*)$. A 'yes' there means no approvable, merger-viable cap would have induced entry, not merely that the proposed one failed. Δ_{remedy} and $\Delta\Pi^M$ vary across rows because the time trend and market size move the no-entry and no-merger equilibria.