Measuring Merger Effects with Revenue Data*

Jacob T. Hess†
Bank of Spain

Xufeng Wang‡
KU Leuven

This Version: June 1, 2026
First Version: February 28, 2025

Abstract

We study how revenue productivity and markups evolve after mergers across the economy. In financial-statement data, targets cease to exist as reporting entities after a deal, so tracking the merged firm requires constructing the combined acquirer-target unit before the merger. We formalize the biases that arise without this correction and implement it across 3,759 horizontal mergers in Spain between 1997 and 2022. Revenue productivity rises by approximately 4 percent within eight years and markups by approximately 5 to 6 percent. The industry decomposition reveals that revenue productivity gains are broad-based, while the overall markup increase is driven by large increases in a few industries, particularly human health activities. In the majority of industries, markup increases are small or negative. We find that without the boundary correction, both estimated effects reverse sign. The results suggest that mergers tend to raise the revenue productivity of the combined firm broadly, consistent with operational improvements, while markup increases are a feature of specific industries rather than a general consequence of consolidation.

Keywords: Mergers and Acquisitions, Revenue Productivity (TFPR), Markups, Consolidation Bias, Virtual Consolidation, Difference-in-Differences

JEL Classification: C23, D24, G34, L41

*We thank Marcela Eslava, Basile Grassi, Pau Roldan-Blanco and Raül Santaeulàlia-Llopis for helpful comments, as well as numerous seminar participants. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Spain or the Eurosystem.

†Email: jacob.hess@bde.es
‡Email: xufeng.wang@kuleuven.be

1 Introduction

Mergers and acquisitions are among the most common forms of corporate restructuring. Whether they improve productive efficiency or increase market power is a central question in industrial organization and competition policy. Yet the empirical evidence remains limited, and the available findings conflict. In the few industries where researchers have been able to measure post-merger outcomes credibly, efficiency gains appear in some settings but not others. Acquired U.S. power plants become roughly 2 percent more efficient within the first few years of new ownership (Demirer & Karaduman, 2024), but plant-level acquisitions in U.S. manufacturing show no corresponding productivity improvement (Blonigen & Pierce, 2016). Markups and prices, by contrast, tend to rise across most settings studied, though the magnitudes and mechanisms vary widely—from coordinated pricing effects of 6 to 8 percent in U.S. brewing (Miller & Weinberg, 2017) to markup increases that accompany quality improvements in Indian manufacturing (Stiebale & Vencappa, 2018). These findings are difficult to reconcile, not least because the studies are not designed to be compared. They cover industries selected for the availability of unusually clean data, and they examine different objects—generators, plants, brand-regions, firm-product pairs. Whether any of these results characterizes what mergers do to the firms involved is an open question. No existing study tracks how the merged firm's productivity and markups evolve broadly across industries.

We study how the merged firm's revenue productivity and markups evolve in the years following a deal, drawing on a representative panel that spans the full range of industries in the economy. Studying mergers at this scale requires financial-statement data, the only data source with broad cross-industry coverage. But in financial-statement data, targets typically cease to exist as separate reporting entities after an acquisition. Their accounts are absorbed into the acquirer, so that a before-after comparison of the acquiring firm compares two different objects: the acquirer alone before the deal and the combined entity afterward. This is a change in what is being measured, not in performance. To maintain a constant unit of observation, we construct the combined acquirer-target entity before the deal by pooling their financial statements, so that the same object, the merged firm, is measured throughout the event window. We implement this approach using firm-level accounts from Bureau van Dijk's Orbis, linked to transaction records from Orbis M&A, across 3,759 horizontal mergers in Spain between 1997 and 2022. Spain offers exceptionally complete firm-level coverage, with mandatory commercial registry filing and one of the longest stable panels among the countries in Orbis.

Revenue productivity measures how efficiently the merged firm converts inputs into revenue, reflecting both operational performance and pricing. In event-study regressions on this combined unit, gross-output TFPR rises by approximately 4 percent within eight years of the merger, and value-added TFPR by approximately 9 percent. Markups, measured by the inverse materials share, rise by approximately 5 to 6 percent over the same horizon; a production-function markup following De Loecker and Warzynski (2012) shows a similar pattern, though the estimates are less precise. The estimates are consistent with parallel trends before the merger. These effects build gradually over several years, a pattern more consistent with operational integration within the merged firm than with immediate financial restructuring. The results are robust across a range of alternative specifications, including a matched difference-in-differences design that controls for selection into mergers.

Whether the gains primarily reflect improved efficiency or increased pricing power, however, varies across industries. Revenue productivity increases are broad-based, with most industries that have significant merger activity showing rising TFPR in the years following a deal. The overall markup increase does not follow the same pattern. It is driven by large increases in a few industries, particularly human health activities, where markups rise sharply with no corresponding productivity gain. The largest merger industries, wholesale and retail trade, show strong TFPR gains alongside small markup effects. The industry decomposition suggests that the conflicting findings in the existing literature may reflect genuine differences across settings: revenue productivity tends to rise following mergers broadly across industries, but whether markups rise depends on the industry.

Prior studies of post-merger outcomes have worked with data where individual production units can be tracked through ownership changes, such as plant-level Census records, generator-level data, and firm-product panels. Such data exist only in the few industries where administrative records follow production units over time. For the rest of the economy, financial-statement data is the only source with broad coverage, and it is precisely the data environment where the measurement problem described above arises.

We formalize the biases that this measurement problem introduces. Without the boundary correction, comparing the acquirer before the deal to the merged entity afterward confounds the real change in performance with a mechanical wedge that reflects the change in the reporting unit, not in the firm's operations. Our production-function estimates indicate near-constant returns to scale across industries. Under this restriction, the wedge in revenue productivity is downward whenever the acquirer is at least as productive as the target. Because the merged firm's accounts pool the acquirer and target, the combined unit's measured productivity is a blend of the two firms' productivities. When the acquirer is the more productive of the two, this blend comes out below the acquirer's standalone level, even though neither firm's operations have changed.

For markups, two distinct biases operate. The first is a boundary bias. After the deal, the consolidated accounts report a single markup for the merged entity, computed from its total revenue and total materials. When the target has a lower markup than the acquirer, adding the target to these accounts pulls the merged entity's measured markup below the acquirer's standalone level, a decline that reflects the change in the reporting unit rather than a genuine fall in the markup. The second bias concerns how the pre-merger combined markup is constructed. This choice arises only before the deal, when the two firms report separately; afterward there is a single entity and a single markup, with nothing to weight. A researcher who has recognized the boundary problem must still decide how to combine the two pre-merger markups into a single object comparable to the merged entity's. Summing the firms' revenue and materials weights the two markups by cost; averaging the separate markups with sales weights instead overstates the pre-merger level, biasing the estimated markup change downward. Pooling accounts before the deal addresses all three biases at once: a single summation holds the measurement boundary constant across the event window and reconstructs the object the post-merger accounts report.

When the same event study is estimated on the raw acquirer panel, where the boundary shifts at the deal date, revenue productivity appears to decline after mergers. Markups show the same reversal. A researcher tracking the acquirer through the deal without correcting for the boundary change would conclude that mergers reduce both the productivity and the markups of the firms involved. The boundary-consistent estimates tell the opposite story. Both revenue productivity and markups rise. The bias is not a second-order adjustment. It is large enough to change the sign of both estimated effects, reversing the qualitative conclusion about what mergers do to the merged firm. The choice of aggregation weights is a separate matter and just as consequential. Even on the boundary-consistent unit, aggregating pre-merger markups with sales weights rather than cost weights is enough to reverse the estimated markup increase. These problems are not specific to our setting. They arise in any study that tracks an acquirer through a merger in firm-level accounting data, and the correction applies wherever the unit of observation changes at the deal date.

The corrected estimates reveal a pattern that is only visible across the full range of industries. In the overall results, both revenue productivity and markups rise after mergers. The industry decomposition, however, reveals that these two increases are not alike. Revenue productivity gains are broad-based, appearing across most industries with significant merger activity, including wholesale and retail trade, food manufacturing, professional services, and construction. The overall increase in markups is not the result of moderate increases across many industries but of large increases in a few, particularly human health activities, where markups rise sharply after mergers. The majority of merger activity occurs in industries where markup increases are small or negative.

This cross-industry variation puts the existing literature's disagreement in a new light. The industry-specific literature has documented mergers that raise markups without improving productivity, mergers that improve efficiency with little effect on pricing, and mergers that do both. The representative-economy benchmark suggests that these are not conflicting answers to the same question but accurate descriptions of different parts of the economy. The disagreement in the existing literature reflects genuine variation across industries.

These studies can often separate efficiency from pricing directly because they observe physical quantities. At economy-wide scale, only revenue data are available. Revenue data cannot separate physical efficiency from pricing, as that decomposition requires quantity data that do not exist across the full range of industries. But the cross-industry variation in our estimates makes revenue productivity informative about the underlying sources of the gains. If TFPR increases were driven entirely by higher prices, they should track markup increases across industries, since a pure price increase raises both measures by the same proportion. In the data, TFPR gains are broad-based while markup increases are concentrated. This asymmetry is difficult to reconcile with price increases as the sole driver of the revenue productivity gains.

In industries where revenue productivity gains substantially exceed markup effects, which together account for the majority of merger activity, the pattern is more consistent with operational improvements within the merged firm. In industries where markups rise sharply without corresponding productivity gains, particularly human health activities, the pattern is more consistent with increased market power. Neither interpretation alone describes the full picture. Distinguishing these two patterns requires the cross-industry decomposition that representative-economy coverage provides.

On balance, the representative-economy evidence suggests that mergers tend to raise the revenue productivity of the combined firm broadly across industries, while markup increases are a feature of specific industries rather than a general consequence of consolidation. These findings provide a benchmark for the post-merger firm across the representative economy, against which future industry-specific evidence can be compared and contextualized.

Related Literature. Relatively few papers study how mergers affect productivity and market power (Asker & Nocke, 2021), and the available evidence comes almost entirely from individual industries. Where physical output can be measured, some studies find efficiency gains at acquired production units (Braguinsky, Ohyama, Okazaki, & Syverson, 2015; Demirer & Karaduman, 2024), while others find that market power effects dominate (Kulick, 2017; Miller & Weinberg, 2017). Ashenfelter, Hosken, and Weinberg (2015) document both forces in the same setting, finding that merger-specific cost reductions in U.S. brewing roughly offset the price increases induced by higher concentration. Stiebale and Vencappa (2018) find that acquisitions in India raise both markups and product quality, exploiting firm-product-level data that allow a full decomposition of efficiency, markups, and quality. Bhattacharya, Illanes, and Stillerman (2023) provide the most comprehensive analysis of merger price effects to date and document substantial heterogeneity across consumer-goods mergers, a finding that resonates with the cross-industry variation we document. A common feature of these studies is that the data allow researchers to track individual production units, such as plants, generators, or products, through ownership changes, avoiding the measurement problem that arises in financial-statement data. This sidesteps the boundary issue but restricts scope to the few industries where such tracking is possible.

Among existing work, two papers are particularly close to ours in approach. Blonigen and Pierce (2016) study the same objects we do, revenue productivity and markups recovered via production-function methods, but use plant-level Census data restricted to U.S. manufacturing. Because plants retain their identifiers through ownership changes, the boundary problem that arises in financial-statement data does not apply in their setting. They find that markups rise with little change in productivity, a pattern that our industry decomposition suggests is specific to certain industries rather than a general feature of mergers. Stiebale and Szücs (2022) use the same data source (Orbis) and recover markups in European markets, but they study the response of rival firms to mergers notified to the European Commission rather than the merged entity itself. We study the merged firm directly, across the representative population of mergers rather than the subset selected for regulatory review.

Outline. Section 2 describes the data, deal linkage, and sample construction. Section 3 formalizes the measurement biases that arise in financial-statement data and describes the virtual consolidation correction. Section 4 presents the event-study evidence on revenue productivity and markups, the industry decomposition, and the consolidation-bias demonstration. Section 5 provides robustness checks. Section 6 concludes.

2 Data

We use firm-level financial statements for Spain from Bureau van Dijk's Orbis and M&A transaction records from Orbis M&A. The financial data are drawn from annual accounts filed with the Spanish Commercial Registry.[1] Among the countries covered in Orbis, Spain offers one of the longest and most complete firm-level panels, with stable coverage of the corporate sector spanning more than two decades. Following Kalemli-Özcan, Sørensen, Villegas-Sanchez, Volosovych, and Yeşiltaş (2024), we construct a representative sample covering 1997–2022, with coverage averaging 63% of registered firms per year; Appendix A.2 compares the sample to Spanish National Accounts. Orbis M&A records more than 8,000 domestic transactions in Spain over this period. We link these to the firm panel using Bureau van Dijk identifiers and the baseline event year is the deal announcement date.[2]

Because Orbis reports consolidated accounts, the surviving entity's accounts absorb the target's activity after an acquisition, and targets typically disappear from the panel as separate reporting units. In the acquirer's panel, the pre-deal observations therefore reflect one firm's operations and the post-deal observations reflect the combined operations of both firms. This shift in the accounting boundary can generate spurious movements in measured productivity and markups even absent any real economic change; we formalize the bias and its direction in Section 3. To ensure that the unit of observation is constant throughout the event window, we construct the analysis at the level of acquirer-target pairs, aggregating the financial statements of both firms in each pre-deal year.

We limit the sample to domestic transactions, since tracking both the acquirer and the target in the firm panel requires both to file with the Spanish registry. We further restrict attention to firms involved in a single transaction during the sample period, because overlapping deals contaminate the event-study design when the post-period of one transaction falls within the pre-period of another. Fewer than 20% of acquirers and fewer than 5% of targets appear in more than one deal. For the productivity and markup analysis, we focus on horizontal transactions, defined as deals in which both parties operate within the same NACE 2-digit industry, the level at which production functions are estimated. More than half of all transactions involve parties in different 2-digit industries; these cross-industry deals are excluded from the baseline but included in a robustness check (Section 5). The resulting sample contains more than 3,700 horizontal mergers spanning the full range of industries in the Spanish economy.

[1] Filing is mandatory under Spanish commercial law; non-compliance is subject to penalties.

[2] The recorded announcement date can lag the target's disappearance from the panel by one or more years. We reconcile event timing with the target's last observation year; Appendix A details the procedure.

Table 1: Pre-Deal Firm Characteristics

	Non-M&A	Acquirers	Targets
Number of Employees	13	220	89
Total Assets	1,860	56,702	19,305
Gross Output	1,891	49,876	15,311
Value Added	558	16,588	4,938
Material Costs	1,334	33,288	10,373
Firms	1,044,359	4,599	4,780

Note: Monetary values are in thousands of constant 2015 euros. Employment is in levels. For acquirers and targets, the table reports means in the year prior to the deal. For non-M&A firms, the table reports pooled averages across the full sample period. Acquirers and targets include only firms involved in exactly one transaction during the sample period.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

Firms that participate in M&A are among the largest in the Spanish economy. Table 1 reports mean characteristics of acquirers and targets in the year before the deal, alongside pooled averages for firms that never participate in M&A.[3] Acquirers average 220 employees and EUR 57 million in total assets, compared to 13 employees and EUR 1.9 million for the typical non-M&A firm (Appendix A.3 compares both groups to the national firm size distribution). Targets are smaller than acquirers but still far larger than non-M&A firms, with 89 employees and EUR 19 million in total assets. Across all size measures, acquirers are roughly two to three times larger than targets.

3 Measurement

3.1 The measurement problem and virtual consolidation

We study how revenue productivity and markups evolve around mergers and acquisitions. This requires comparing outcomes for the merged entity before and after the deal on a consistent unit of observation. Post-merger financial statements, however, change the unit being measured. When a target is absorbed, its activity is folded into the acquirer's consolidated accounts and it ceases to exist as a separate reporting entity. The observed post-merger accounts therefore correspond to pooled outcomes for the combined entity.

$$Q_t^{\text{pooled}} \equiv Q_{A,t} + Q_{T,t}, \qquad K_t^{\text{pooled}} \equiv K_{A,t} + K_{T,t},$$
$$L_t^{\text{pooled}} \equiv L_{A,t} + L_{T,t}, \qquad M_t^{\text{pooled}} \equiv M_{A,t} + M_{T,t},$$

Here A denotes the acquirer, T the target, and Q, K, L, M denote gross output, capital, labor, and intermediate inputs (materials expenditure), respectively.

Let f denote the production function, assumed to be strictly increasing, concave, and differentiable.

[3] Results are similar when restricting the non-M&A sample to the same calendar-year distribution as the pre-deal observations.

Revenue-based total factor productivity (TFPR) for an individual firm is

$$Z_{it} \equiv \frac{Q_{it}}{f(K_{it}, L_{it}, M_{it})}, \tag{1}$$

and for the pooled entity,

$$Z_t^{\text{pooled}} \equiv \frac{Q_t^{\text{pooled}}}{f\left(K_t^{\text{pooled}},\, L_t^{\text{pooled}},\, M_t^{\text{pooled}}\right)}. \tag{2}$$

A naive comparison of Z_A before the deal with Z^{pooled} afterward confounds any real change in performance with the mechanical change in the reporting boundary. The pre-merger object measures productivity for the acquirer alone, while the post-merger object measures productivity for the combined entity. Even if neither firm's technology, prices, or input usage changed, Z^{pooled} is generically different from Z_A because it evaluates the production function at a different input bundle and scales it by a different level of output.

The same problem applies to markups. Following De Loecker and Warzynski (2012), markups are measured using the first-order condition for a variable input. For materials, the firm-level markup is

$$\mu_{it} = \frac{\theta^M(\mathbf{x}_{it})}{\alpha_{it}^M}, \tag{3}$$

where $\theta^M(\mathbf{x}_{it}) \equiv \frac{M_{it} f_M(\mathbf{x}_{it})}{f(\mathbf{x}_{it})}$ is the output elasticity of materials evaluated at the firm's input bundle $\mathbf{x}_{it} \equiv (K_{it}, L_{it}, M_{it})$, $f_M \equiv \partial f/\partial M$ denotes the marginal product of materials, and $\alpha_{it}^M \equiv M_{it}/Q_{it}$ is the materials expenditure share. A comparison of the acquirer's pre-merger markup μ_A with the merged entity's post-merger markup μ^{pooled} inherits the same boundary mismatch.

To restore comparability, we construct for each deal a virtual consolidated pre-merger entity by summing the acquirer's and target's accounts prior to the transaction in the same way that post-merger consolidation combines them. The unit of observation is therefore the acquirer–target pair $A{+}T$ throughout the event window, and our estimands are the event-time paths of $\ln Z_t^{\text{pooled}}$ and $\ln \mu_t^{\text{pooled}}$, both evaluated on a constant accounting boundary before and after the deal. Because these objects are computed from total output, capital, labor, and materials at the level of the combined entity, they are invariant to internal bookkeeping (reallocation of overhead, routing of intermediate inputs, or booking of revenues across subsidiaries). Appendix A.3.4 details the pair construction and the associated sample definitions.

3.2 Production function specification and estimation

Computing pooled TFPR and markups requires specifying the production function f, and the choice of functional form matters for the before–after comparison. From (3), the change in the pooled markup decomposes as

$$\ln \mu_1^{\text{pooled}} - \ln \mu_0^{\text{pooled}} = \underbrace{\left[\ln \theta^M(\mathbf{x}_1^{\text{pooled}}) - \ln \theta^M(\mathbf{x}_0^{\text{pooled}})\right]}_{\text{elasticity channel}} + \underbrace{\left[\ln \alpha_0^{M,\text{pooled}} - \ln \alpha_1^{M,\text{pooled}}\right]}_{\text{cost-share channel}}, \quad (4)$$

where $\theta^M(\mathbf{x}_j^{\text{pooled}})$ is the output elasticity of materials evaluated at the pooled input bundle in period j. The cost-share channel is where real economic forces operate (efficiency, market power, input prices). An analogous decomposition holds for TFPR. Under a general f, the change in $\ln f(\mathbf{x}^{\text{pooled}})$ depends on how the output elasticities vary as the pooled input bundle shifts from pre- to post-merger, and this variation contaminates the measured productivity change.[4] In both decompositions, the elasticity channel captures the mechanical drift of the output elasticities as the pooled input bundle shifts between pre- and post-merger periods. When the merger reshuffles inputs across the combined entity, the point at which elasticities are evaluated changes, and this alone can move measured TFPR and markups even if neither firm's technology, prices, or input usage changed.

Under a flexible specification such as the translog, output elasticities are functions of the input bundle. For example, the materials elasticity takes the form

$$\theta_{it}^M = \beta_m + \beta_{mm} \ln M_{it} + \beta_{mk} \ln K_{it} + \beta_{ml} \ln L_{it} + \cdots,$$

so when the merger reshuffles inputs across the combined entity, θ^M drifts between pre- and post-merger periods. The elasticity channel is generically nonzero, and the researcher cannot separate it from the cost-share channel without knowing the counterfactual input allocation that would have prevailed absent the merger.

We adopt a Cobb–Douglas gross-output specification,

$$Q_{it} = e^{\omega_{it} + \epsilon_{it}} K_{it}^{\beta_k} L_{it}^{\beta_l} M_{it}^{\beta_m}, \quad (5)$$

where ω_{it} is firm productivity anticipated by the firm when choosing inputs and ϵ_{it} captures unanticipated shocks and measurement error. Under Cobb–Douglas, output elasticities are constants $(\beta_k, \beta_l, \beta_m)$ that do not depend on the input bundle. The elasticity channel in (4) is therefore identically zero, and the same holds for the TFPR parallel. Every observed change in pooled TFPR or pooled markups reflects actual changes in output relative to inputs, not the output elasticities drifting with the input mix. Pooled TFPR reduces to

$$Z_t^{\text{pooled}} = \frac{Q_t^{\text{pooled}}}{(K_t^{\text{pooled}})^{\beta_k} (L_t^{\text{pooled}})^{\beta_l} (M_t^{\text{pooled}})^{\beta_m}}, \quad (6)$$

[4] Taking a first-order expansion of $\ln f$ around the pre-merger pooled inputs, $\Delta \ln Z^{\text{pooled}} = \left[\Delta \ln Q^{\text{pooled}} - \sum_{j \in \{K,L,M\}} \theta_0^j \, \Delta \ln x_j^{\text{pooled}}\right] - R(\mathbf{x}_0^{\text{pooled}}, \mathbf{x}_1^{\text{pooled}})$, where θ_0^j is the output elasticity of input j at the pre-merger pooled bundle and $R = \int_0^1 \sum_j \left[\theta^j(\mathbf{x}(s)) - \theta_0^j\right] \frac{d \ln x_j(s)}{ds} \, ds$ captures the drift of elasticities along the path from $\mathbf{x}_0^{\text{pooled}}$ to $\mathbf{x}_1^{\text{pooled}}$. Under Cobb–Douglas, $R \equiv 0$.

and the pooled markup to

$$\mu_t^{\text{pooled}} = \frac{\beta_m}{\alpha_t^{M,\text{pooled}}} = \beta_m \cdot \frac{Q_t^{\text{pooled}}}{M_t^{\text{pooled}}}, \tag{7}$$

both computed from observables and fixed parameters.

Estimation. We estimate the output elasticities $(\beta_k, \beta_l, \beta_m)$ by industry (NACE 2-digit) using the control-function approach of Ackerberg, Caves, and Frazer (2015). In logs, the production function is

$$q_{it} = \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \omega_{it} + \epsilon_{it},$$

where ω_{it} is the productivity component anticipated by the firm when choosing inputs but unobserved by the econometrician. Materials are treated as a flexible input whose demand is strictly increasing in ω_{it} conditional on predetermined capital, dynamic labor, and observed demand shifters (firm location, firm type, listing status, and year fixed effects). Under this monotonicity condition, the materials demand function can be inverted to proxy for ω_{it}, and the output elasticities are identified from GMM moment conditions exploiting the assumption that productivity follows a first-order Markov process. Within this framework, identifying the output elasticity of a flexible input is a well-known challenge. Gandhi, Navarro, and Rivers (2020) show that under perfect competition, a firm's materials demand is fully determined by its productivity and predetermined inputs, leaving lagged materials without the independent variation needed to pin down the materials coefficient. Following Ackerberg and De Loecker (2024), we use lagged materials as an instrument for current materials. Strategic interactions among producers generate variation in firms' input choices beyond each firm's own productivity, providing the identifying variation needed for the materials elasticity. Appendix B reports the full estimation procedure and the estimated elasticities by industry.

We restrict the estimation sample to firms that are not involved in M&A during the study period. M&A transactions can induce abrupt changes in the accounting unit and can alter market power and input choices in ways that threaten the monotonicity condition used for identification. We estimate the production function on non-M&A firms and then apply the estimated elasticities to construct TFPR and markups for the acquirer–target pairs analyzed in the event study. As a robustness check, Appendix B.2 re-estimates the production function including M&A-involved firms in the sample and applying the selection-correction procedure of De Loecker, Goldberg, Khandelwal, and Pavcnik (2016) to account for the endogeneity of M&A participation.

Cross-industry deals. The baseline focuses on horizontal mergers, where the acquirer and target operate within the same NACE 2-digit industry and share common output elasticities. In cross-industry deals, the acquirer and target generally possess different elasticities $(\beta_A \neq \beta_T)$, and the pooled inputs of two distinct technologies do not naturally map to output via a single production function. As a robustness exercise, we compute pooled TFPR and markups for cross-industry deals using the acquirer's industry elasticities in both the pre- and post-merger periods. This ensures that the same parameter vector is applied to the same accounting boundary before and after the deal, but it imposes the acquirer's technology on the target's inputs and therefore lacks the structural interpretation of the horizontal case. Appendix F.4 reports the results.

3.3 Characterizing Measurement Bias

The Cobb–Douglas specification delivers closed-form expressions for the measurement biases that virtual consolidation corrects. These expressions pin down the conditions under which each bias is positive or negative at the deal level. Formal results and proofs are in Appendix G.

3.3.1 Acquirer-only TFPR bias

Without virtual consolidation, the researcher compares the acquirer's pre-merger TFPR with the pooled TFPR of the merged entity after the deal. The acquirer-only estimate can be written as

$$\underbrace{\ln Z^{\text{pooled}}_{\text{post}} - \ln Z_{A,\text{pre}}}_{\text{acquirer-only estimate}} = \underbrace{\ln Z^{\text{pooled}}_{\text{post}} - \ln Z^{\text{pooled}}_{\text{pre}}}_{\text{boundary-consistent estimate}} + \underbrace{\ln Z^{\text{pooled}}_{\text{pre}} - \ln Z_{A,\text{pre}}}_{\text{consolidation bias}}, \tag{8}$$

where $Z^{\text{pooled}}_{\text{pre}}$ is TFPR for the virtual consolidated pre-merger entity constructed in Section 3.1. The first term compares the same object before and after the deal. The second is the pre-merger gap between pooled and acquirer TFPR, a deal-specific mechanical wedge that has nothing to do with performance. The Cobb–Douglas specification provides an exact characterization of this wedge.

Proposition 1 (Consolidation bias identity and threshold)**.** *Consider two production units $i \in \{A, T\}$ operating the Cobb–Douglas technology (5) with common output elasticities $(\beta_k, \beta_l, \beta_m)$. Let*

$$s_K \equiv \frac{K_A}{K^{\text{pooled}}}, \quad s_L \equiv \frac{L_A}{L^{\text{pooled}}}, \quad s_M \equiv \frac{M_A}{M^{\text{pooled}}}$$

denote the acquirer's shares of pooled capital, labor, and materials, let $r \equiv Z_T/Z_A$ denote the target's relative TFPR, and define

$$a \equiv s_K^{\beta_k}\, s_L^{\beta_l}\, s_M^{\beta_m}, \qquad b \equiv (1-s_K)^{\beta_k}\,(1-s_L)^{\beta_l}\,(1-s_M)^{\beta_m}.$$

Then the consolidated-to-acquirer TFPR ratio satisfies

$$\frac{Z^{\text{pooled}}}{Z_A} = a + r\,b,$$

and $Z^{\text{pooled}} \gtrless Z_A$ if and only if $r \gtrless r^* \equiv (1-a)/b$.

The proof is in Appendix G. Proposition 1 provides an analytical expression for the consolidation bias under Cobb–Douglas. The direction of the bias is ambiguous and depends on the characteristics of the firms involved in the deal. The bias for each deal is fully determined by two observable quantities: the within-pair input allocation and the acquirer–target productivity gap. Specifically, the consolidation bias equals

$$\ln Z_{\text{pre}}^{\text{pooled}} - \ln Z_{A,\text{pre}} = \ln(a + r\, b),$$

where a, b, and r are evaluated at pre-merger values. The weights a and b are the Cobb–Douglas production function evaluated at the acquirer's and target's respective shares of pooled inputs, capturing each firm's contribution to the combined entity's production capacity. The ratio r measures the target's productivity relative to the acquirer's. For a given deal, the bias is negative when $r < r^*$ and positive when $r > r^*$.

Our production function estimates indicate near-constant returns to scale in nearly all industries (Table B.1). Under this restriction, the consolidation bias is weakly negative for all deals where the acquirer is at least as productive as the target, and weakly negative even when productivity is equal. Recall from (6) that pooled TFPR is total revenue divided by the production function evaluated at combined inputs. Under constant returns, concavity implies

$$f(K^{\text{pooled}}, L^{\text{pooled}}, M^{\text{pooled}}) \;\geq\; f(K_A, L_A, M_A) \;+\; f(K_T, L_T, M_T),$$

so the denominator of pooled TFPR exceeds the sum of the individual firms' denominators. Total revenue, by contrast, adds linearly. Pooled TFPR therefore falls below the acquirer's standalone TFPR, making the consolidation bias in (8) negative.

Under constant returns, the consolidation bias can remain negative even when the target is more productive than the acquirer. Concavity implies $a + b \leq 1$, so the threshold satisfies $r^* \geq 1$ (Corollary 1, Appendix G). The direction of the bias then depends on the acquirer–target productivity gap. When the acquirer is more productive than the target ($r < 1$), we have $r < 1 \leq r^*$, so the bias is strictly negative. When both firms are equally productive ($r = 1$), the bias equals $\ln(a + b) \leq 0$. When the target is more productive ($r > 1$), the bias is negative if r remains below r^* and positive if r exceeds it. Additional special cases, including constant input ratios, are in Corollary 2 (Appendix G).

The event study aggregates across deals, so the aggregate bias in the acquirer-only specification is the average of the deal-specific terms $\ln(a + r\, b)$. The aggregate direction depends on the distribution of relative productivities and input allocations across deals in the sample. We compare the boundary-consistent and acquirer-only specifications in Section 4 to assess the empirical relevance and direction of the aggregate bias.

3.3.2 Acquirer-only markup bias

The same boundary problem arises for markups, but the structure is simpler. Without any correction, the researcher compares the acquirer's pre-merger markup with the pooled post-merger markup of the combined acquirer–target entity. These two objects are not comparable. The post-merger markup includes the target, so the estimated change can reflect the mechanical addition of the target rather than a change in the acquirer's pricing behavior.

Because the Cobb–Douglas markup is the ratio $\beta_m Q_i / M_i$, and both revenue and materials add linearly when accounts are pooled, the merged entity's markup can be decomposed as

$$\mu^{\text{pooled}} = \beta_m \frac{Q_A + Q_T}{M_A + M_T} = \frac{M_A}{M^{\text{pooled}}} \cdot \frac{\beta_m Q_A}{M_A} + \frac{M_T}{M^{\text{pooled}}} \cdot \frac{\beta_m Q_T}{M_T} = w_A \mu_A + w_T \mu_T,$$

where $w_i \equiv M_i / M^{\text{pooled}}$. The M_i terms cancel in each product, so the pooled markup from consolidated accounts is exactly the cost-weighted average of the individual markups. Because a weighted average lies between its components, the pooled markup falls between the acquirer's and target's pre-merger markups. When the target has a lower markup than the acquirer, the pooled figure sits below the acquirer's standalone level, and comparing the acquirer before the deal with the merged entity afterward shows a decline that is purely mechanical. The mechanical wedge from comparing the acquirer with the pooled entity is

$$\mu^{\text{pooled}} - \mu_A = w_T \left(\mu_T - \mu_A\right).$$

The direction depends only on which firm has the higher markup, unlike the TFPR consolidation bias, where input shares, productivity ratios, and returns to scale all interact (Proposition 1).

The event study aggregates across deals, so the aggregate direction of the markup boundary bias depends on the distribution of relative markups across deals in the sample. We compare the boundary-consistent and acquirer-only specifications in Section 4 to assess its empirical relevance and direction. Virtual consolidation addresses both the TFPR and markup boundary biases by holding the measurement boundary constant throughout the event window.

3.3.3 Pre-merger markup aggregation

The construction of the pre-merger combined markup also affects the estimated effect. The literature commonly uses either sales-weighted or cost-weighted averages. In our setting, the post-merger markup is observed from pooled consolidated accounts, and the markup computed from pooled accounts is equivalent to a materials cost-weighted average of the underlying firm markups. Our acquirer-target virtual-consolidation measure therefore has a direct interpretation. It compares the cost-weighted markup of the combined acquirer–target unit before the merger with the cost-weighted markup of the same unit after the merger. If instead the pre-merger combined markup is constructed with sales weights, the comparison changes. The researcher is then comparing a sales-weighted pre-merger markup with a cost-weighted post-merger markup, because the post-merger pooled accounts imply cost weighting. This aggregation mismatch can generate a mechanical downward bias in measured markups even when the merger has no real effect.

Proposition 2 (Sales-weighted aggregation exceeds the pooled markup)**.** *Maintain the assumptions of Proposition 1. For a given deal and year, let* $\pi_i \equiv Q_i / (Q_A + Q_T)$ *denote sales weights and* *define the sales-weighted markup* $\mu^{\text{sales}} \equiv \pi_A \, \mu_A + \pi_T \, \mu_T$*. Then*

$$\mu^{\text{sales}} \;\geq\; \mu^{\text{pooled}},$$

with equality if and only if $\alpha_A^M = \alpha_T^M$ *(the materials expenditure shares are identical across the acquirer and target).*

The proof is in Appendix G. The result follows from the convexity of the markup function in the cost share. Under the Cobb–Douglas formula, the markup is β_m divided by the materials expenditure share $c_{it} \equiv M_{it}/Q_{it}$, and $f(c) = \beta_m/c$ is strictly convex. The pooled cost share equals the sales-weighted average of the individual cost shares, since

$$\frac{M^{\text{pooled}}}{Q^{\text{pooled}}} \;=\; \pi_A \, \frac{M_A}{Q_A} + \pi_T \, \frac{M_T}{Q_T} \;=\; \pi_A \, c_A + \pi_T \, c_T,$$

and the pooled markup equals f evaluated at this average. By Jensen's inequality, the sales-weighted average of individual markups, $\pi_A \, f(c_A) + \pi_T \, f(c_T)$, exceeds $f(\pi_A \, c_A + \pi_T \, c_T) = \mu^{\text{pooled}}$, with equality only when $c_A = c_T$.

Proposition 2 shows why the sales-weighted markup can generate a mechanical downward bias. Suppose that the underlying firm-level markups and cost shares do not change before and after the merger. If the pre-merger markup is computed as $\mu_{\text{pre}}^{\text{sales}}$ but the post-merger markup is observed from pooled accounts as $\mu_{\text{post}}^{\text{pooled}}$, then the measured change is

$$\mu_{\text{post}}^{\text{pooled}} - \mu_{\text{pre}}^{\text{sales}}.$$

When nothing changes in the underlying firms, $\mu_{\text{post}}^{\text{pooled}} = \mu_{\text{pre}}^{\text{pooled}}$. Therefore,

$$\mu_{\text{post}}^{\text{pooled}} - \mu_{\text{pre}}^{\text{sales}} = \mu_{\text{pre}}^{\text{pooled}} - \mu_{\text{pre}}^{\text{sales}} \leq 0.$$

Thus a naive comparison can find a negative effect of M&A on markups even when there is no real change in markups at all.

The aggregation bias is distinct from the boundary bias characterized in Section 3.3.1 and 3.3.2: the boundary bias arises from comparing different entities before and after the deal, while the aggregation bias arises from how the pre-merger combined markup is constructed for a given entity. Virtual consolidation addresses both simultaneously, because pooling accounts both holds the measurement boundary constant and constructs the pre-merger markup in the same way the post-merger accounts report it. We compare the two aggregation methods empirically in Section 4.

4 Empirical Results

We study the post-merger dynamics of the boundary-consistent acquirer–target unit across 3,759 horizontal mergers using the staggered event-study estimator of De Chaisemartin and d'Haultfoeuille (2024). Revenue productivity and markups both rise overall, but the industry decomposition reveals that TFPR gains are broad-based while the overall markup increase is driven by large gains in a few industries.

Mergers in the sample span 1997–2022, so acquirer–target pairs enter the post-merger period in different calendar years. Treatment is staggered, and the appropriate estimator must account for this. We adopt an event-study design, with event time measured relative to the announcement year, for two reasons. First, the event-study framework provides a direct check on the parallel-trends assumption by tracing the dynamics of each outcome in the years preceding the deal. If the boundary-consistent unit exhibited differential trends before the merger, the credibility of the post-merger estimates would be undermined. Second, the event-study framework reveals the trajectory of post-merger adjustment rather than collapsing it into a single average. Whether productivity and markup gains materialize immediately or build over several years has different implications for the economic mechanisms at work.

The standard two-way fixed effects (TWFE) event-study estimator is not appropriate for this design. Under staggered treatment adoption, TWFE pools all cohorts into a single set of event-time coefficients. When treatment effects vary across cohorts, as is plausible when mergers of different vintages occur in different macroeconomic environments and involve firms of different sizes, the TWFE estimate at a given event-time horizon can be contaminated by effects from other horizons belonging to cohorts that entered treatment earlier or later (De Chaisemartin & d'Haultfoeuille, 2020; Sun & Abraham, 2021). The resulting coefficients need not recover a weighted average of the true cohort-specific effects and can even have the wrong sign. A growing literature addresses this by estimating cohort-specific treatment effects separately and then aggregating them into event-time averages using well-defined weights (Borusyak, Jaravel, & Spiess, 2024; Callaway & Sant'Anna, 2021; De Chaisemartin & d'Haultfoeuille, 2024; Sun & Abraham, 2021).

Let $\ell \equiv t - E_i$ denote event time relative to the announcement year E_i for pair i, with $\ell = -1$ as the omitted reference period. We estimate

$$\mathcal{Y}_{it} = \alpha_i + \gamma_t + \sum_{c \in \mathcal{C}} \sum_{\ell \neq -1} \beta_{c\ell} \, \mathbb{1}\{E_i = c\} \, \mathbb{1}\{t - E_i = \ell\} + \varepsilon_{it}, \tag{9}$$

where α_i and γ_t are pair and calendar-year fixed effects and $\beta_{c\ell}$ captures the treatment effect for announcement-year cohort c at event time ℓ. We aggregate cohort-specific effects into event-time average treatment effects using the De Chaisemartin and d'Haultfoeuille (2024) estimator with cohort-size weights and cluster standard errors at the pair level.[5]

[5]For a balanced panel, the event-time average treatment effect on the treated (ATT) is

$$ATT_\ell = \sum_c \mathbb{1}\{c + \ell \leq T\} \, (N_c / N_\ell) \, \beta_{c\ell},$$

This estimator accommodates both never-involved firms and not-yet-treated pairs in the control group. Results are robust to the alternative estimators of Callaway and Sant'Anna (2021), Sun and Abraham (2021) and Borusyak et al. (2024) (Appendix D.2). Restricting the control group to not-yet-treated pairs only does not affect the findings (Appendix D.1).

4.1 Revenue Productivity

The industry-specific literature suggests weakly positive priors for TFPR, though through different channels depending on the setting. Mergers raise physical productivity by 2–5% in U.S. power plants, where generators are a homogeneous asset and efficiency is measured directly in physical units (Demirer & Karaduman, 2024). In U.S. manufacturing, plant-level productivity does not change but markups rise (Blonigen & Pierce, 2016). In Indian manufacturing, acquisitions raise both markups and product quality while lowering marginal costs (Stiebale & Vencappa, 2018). Because TFPR conflates efficiency and pricing, a positive response is consistent with all three patterns. Whether any of them characterizes the representative economy, however, is an open question, since each study examines a different object (generators, plants, firm-products) in an industry selected for data availability.

Figure 1 reports the overall TFPR event studies. Gross-output TFPR of the merged entity rises over the post-merger horizon, reaching approximately 4% by year eight (Panel A; Table C.1). The increase builds gradually. Point estimates rise from the deal date and continue to build steadily over the post-merger horizon, consistent with restructuring that plays out over several years. Pre-event estimates are close to zero in the years immediately before the merger, supporting the parallel-trends assumption on the boundary-consistent unit. Value-added TFPR displays the same qualitative trajectory at larger magnitudes, reaching approximately 9% by year eight and becoming statistically distinguishable from zero at longer horizons (Panel B). The larger value-added magnitude is expected, since value-added TFPR does not condition on materials and therefore captures changes in the materials share alongside changes in productivity. The markup increase documented in Section 4.2 operates through this channel.

The gradual trajectory is informative about the nature of the post-merger changes. A pattern driven entirely by repricing or accounting reclassification would produce a discrete level shift at $\ell = 0$ with little subsequent movement. The continued buildup over several years points toward real-side restructuring, including organizational integration, rationalization of production across the combined entity, and redeployment of capital. This pattern echoes the plant-level evidence in Demirer and Karaduman (2024), where productivity gains in acquired power generators phase in over several years as new owners reorganize operations, and the firm-level evidence in Fons-Rosen, Kalemli-Ozcan, Sørensen, Villegas-Sanchez, and Volosovych (2021), where TFPR of foreign-acquired European manufacturing firms rises gradually and becomes significant only after four years.

Because TFPR is revenue-based, the post-merger increase is consistent with genuine efficiency

where $\beta_{c\ell}$ is the cohort-specific treatment effect at horizon ℓ, N_c is the number of pairs in cohort c and N_ℓ is the total number of pairs observed at horizon ℓ.

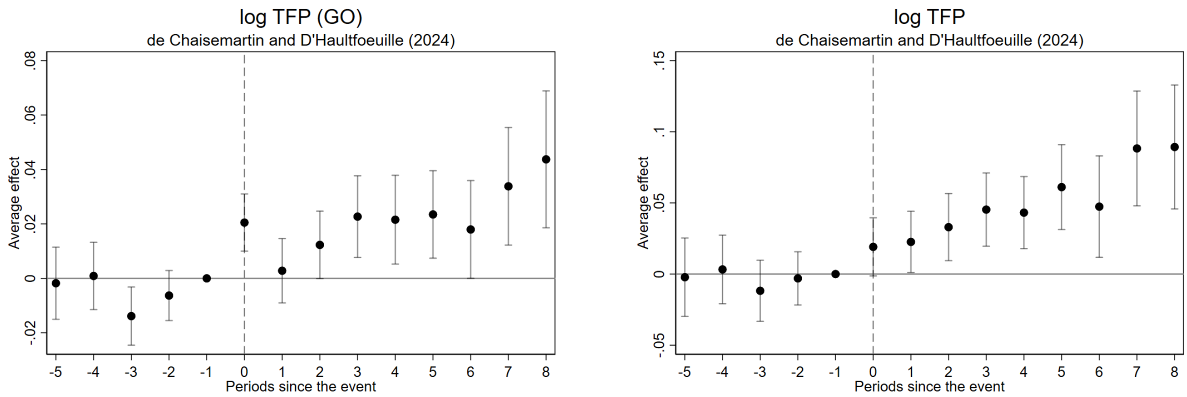

(a) Gross Output

(b) Value Added

Figure 1: Effect of M&A on Revenue Productivity (TFPR)

Note: Event-study estimates of the effect of horizontal M&A on log TFPR for the pooled acquirer–target unit, using the De Chaisemartin and d'Haultfoeuille (2024) estimator. The control group includes firms never involved in a merger and not-yet-treated pairs. Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

gains but equally consistent with the firm raising prices relative to costs. Markups, the ratio of price to marginal cost, speak to this distinction directly.

4.2 Markups

Horizontal mergers raise markups in U.S. manufacturing (Blonigen & Pierce, 2016), after acquisitions in Indian manufacturing (Stiebale & Vencappa, 2018), and among rivals of merging firms in European markets notified to the European Commission (Stiebale & Szücs, 2022). These studies examine specific industries or selected sets of large, policy-relevant deals, so whether the post-merger markup increase characterizes mergers across the broader economy is unknown.

Markups of the merged entity rise after the deal. We measure markups two ways. Our primary measure is the inverse materials share, $\ln(Q_{it}/M_{it})$, which equals the log markup up to an additive constant under Cobb–Douglas (see equation 3), so its event-time variation is identical to the variation in the log markup without requiring production-function estimates. The concern is that measurement error in gross output enters the ratio directly and could move the inverse share mechanically. As a robustness check, we also report a production-function markup that replaces raw gross output with a fitted value from the estimated production function, purging revenue-based noise from the measure.[6]

The inverse materials share rises by approximately 6% by year eight (Figure 2, Panel A; Table C.1), with confidence intervals excluding zero from the deal year onward. The increase is gradual, mirroring the TFPR trajectory, and is present from the first post-event year. Pre-event estimates are close to zero, supporting parallel trends. That markups also build gradually suggests that

[6]The production-function markup is $\widehat{\mu}_{it} = \widehat{\beta}_m \cdot \widetilde{Q}_{it}/M_{it}$, where $\widehat{\beta}_m$ is the estimated materials elasticity and $\widetilde{Q}_{it} \equiv \exp(\widehat{q}_{it})$ is fitted gross output from $q_{it} = f_t(k_{it}, l_{it}, m_{it}, MA_{it}) + \varepsilon_{it}$, a flexible function of inputs and year effects that allows treated observations to shift the mapping via an M&A indicator.

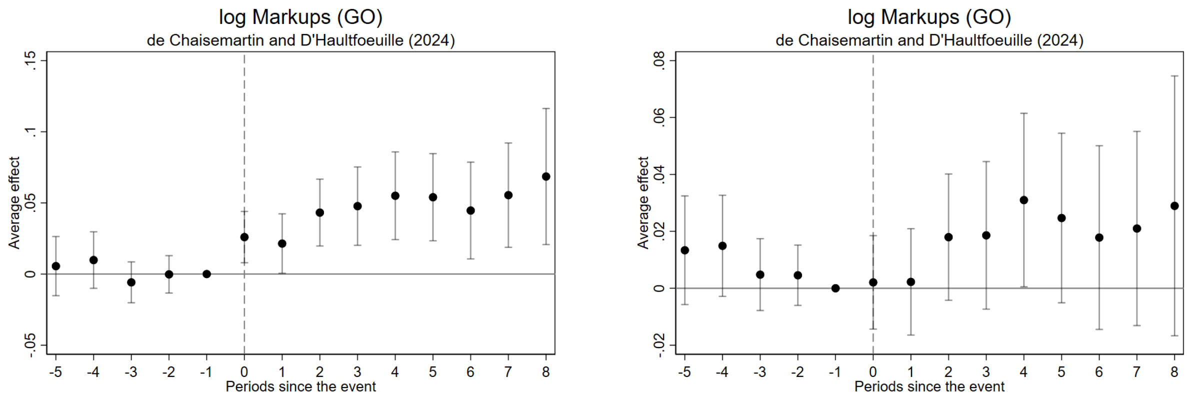

(a) Inverse Materials Share

(b) Production-Function Markup

Figure 2: Effect of M&A on Markups

Note: Event-study estimates of the effect of horizontal M&A on markups for the pooled acquirer–target unit, using the De Chaisemartin and d'Haultfoeuille (2024) estimator. Panel A reports the inverse materials share $\ln(Q_{it}/M_{it})$, which under Cobb–Douglas is proportional to the log markup. Panel B reports the production-function markup using fitted gross output. The control group includes firms never involved in a merger and not-yet-treated pairs. Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

the increase reflects the same slow restructuring process documented for TFPR rather than an immediate repricing of output. The production-function markup, which guards against measurement error, shows the same directional pattern at a more conservative magnitude of approximately 3% by year eight (Panel B). The estimates are less precise. The two measures agree on direction and trajectory, which is more informative than their difference in magnitude; the gap reflects the removal of revenue-side noise from the production-function measure, and the estimates bracket a plausible range for the post-merger markup increase. Unlike the inverse materials share, pre-event estimates for this measure are slightly above zero at longer pre-merger horizons, though they are not significantly different from zero in the three years preceding the merger.

The markup results narrow the interpretation of the TFPR findings. Revenue productivity rises and markups rise; part of the post-merger TFPR gain reflects a widening of the ratio of price to marginal cost, not pure efficiency. But rising markups do not by themselves rule out efficiency gains (Syverson, 2019). The markup is the ratio of price to marginal cost, and it can rise even as both decline. What matters is that price falls less than marginal cost. A merged entity that lowers its costs through restructuring but does not fully pass those savings through to prices would exhibit rising markups alongside genuine efficiency improvements. The TFPR and markup results together establish that the merged entity operates at a higher ratio of price to marginal cost after the deal. Whether this overall pattern characterizes mergers uniformly across industries, or masks important heterogeneity, is what we examine next.

Table 2: Post-Merger Effects in High-Activity Merger Industries

Industry	Share of Mergers (%)	TFPR		Markups	
		Gross Output	Value Added	Inverse Materials Share	Production Function
Wholesale trade (46)	18.3	0.017***	0.067***	0.007	−0.026***
Retail trade (47)	5.0	0.037***	0.101***	0.027***	−0.020***
Motor vehicle trade (45)	4.9	0.029***	0.087***	0.027**	0.013
Land transport (49)	4.4	−0.007	−0.020	−0.090**	−0.110***
Human health activities (86)	3.5	0.000	−0.013	0.161***	0.196***
Food manufacturing (10)	3.4	0.048***	0.100***	0.023*	−0.017
Specialized construction (43)	2.6	0.039*	−0.014	0.124***	0.088**

Note: Average post-merger treatment effects (average of event-time ATTs from $\ell = 0$ to $\ell = 8$) estimated using the De Chaisemartin and d'Haultfoeuille (2024) estimator on the pooled acquirer-target unit. The control group includes firms never involved in a merger and not-yet-treated pairs. Industries are ordered by their share of total merger activity. NACE industry descriptions are as follows. Wholesale trade (46), sale of goods to businesses and retailers. Retail trade (47), sale of goods to consumers. Motor vehicle trade (45), sale, maintenance, and repair of motor vehicles. Land transport (49), freight and passenger transport by road and rail. Human health activities (86), hospitals, medical and dental practices. Food manufacturing (10), processing and preservation of food products. Specialized construction (43), electrical installation, plumbing, and building completion. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

4.3 Heterogeneity Across Industries

The overall estimates in Sections 4.1 and 4.2 combine 3,759 mergers across every industry in the economy. Whether these patterns generalize across industries or whether the overall result obscures important heterogeneity is unclear. We estimate industry-level event studies for the industries with the most merger activity.

Of the top ten industries by deal count, seven satisfy the parallel-trends assumption, collectively accounting for approximately 42% of all mergers in the sample.[7] Two patterns stand out in the industry-level estimates (Table 2). TFPR gains are broad-based. Markup effects are generally small, but the overall increase is driven by large gains in a few industries.

Five of seven interpretable industries show positive and statistically significant post-merger gross-output TFPR, including the three largest merger industries and food manufacturing. Only human health activities and land transport show no significant TFPR gains. For specialized construction, gross-output TFPR is positive but significant only at the 10 percent level, and value-added TFPR is insignificant. The overall TFPR increase reflects a broad pattern, not one driven by a few outlier industries.

Markup dynamics are different and depend on the measure. The inverse materials share rises significantly in several industries, but the largest increases are an order of magnitude above the rest. The production-function markup rises significantly only in two industries. In human health activities, the inverse materials share rises by 16% and the production-function markup by 20%.

[7] Three industries produce unreliable estimates. Construction of buildings (NACE 41, rank 7) exhibits unstable estimates from small sample sizes at longer horizons. Computer programming and consultancy (NACE 62, rank 8) and fabricated metal products (NACE 25, rank 9) have pre-trend violations. We report these in Appendix E.

In specialized construction, the inverse materials share rises by 12% and the production-function markup by 9%.

In the two largest merger industries, wholesale and retail trade, which together account for 23% of all deals, markup effects are small or mixed. In wholesale trade, the inverse materials share shows no significant change while the production-function markup declines significantly. In retail trade, the two measures point in opposite directions: the inverse materials share rises modestly (2.7%) while the production-function markup declines (−2.0%). Food manufacturing, with 3.4% of deals, shows strongly rising TFPR. The inverse materials share rises modestly (2.3%), though not at the 5% significance level, while the production-function markup is statistically insignificant. None of the largest merger industries show markup increases approaching the magnitudes in human health activities and specialized construction. The overall markup increase documented in Section 4.2 is therefore not the result of moderate increases spread across many industries. It is driven by the magnitude of the increases in these two industries, which are large enough to pull the overall result despite the majority of merger activity occurring in industries where markup effects are small or negative.

The industry-level patterns sharpen the interpretation of the overall results. In wholesale trade, retail trade, and food manufacturing, which collectively account for approximately 27% of all mergers, TFPR gains are several times larger than any markup increase. A pure price increase raises TFPR and markups by the same proportion, so most of the revenue productivity improvement in these industries reflects operational changes within the merged firm rather than wider price-cost margins.

By contrast, human health activities presents the opposite picture, with large markup increases and statistically insignificant TFPR. The merged health entity charges substantially more relative to marginal cost without a corresponding improvement in revenue productivity. Land transport completes the spectrum, with statistically insignificant TFPR and declining markups, suggesting mergers in this industry neither improve revenue productivity nor increase pricing power.

The industry decomposition teaches two things. First, the TFPR finding generalizes. Revenue productivity gains appear across most interpretable industries, regardless of whether markups rise or fall. Second, the markup finding does not. The overall markup increase is driven by specific industries, not by a uniform pattern across the economy.

4.4 The Consolidation Bias in Practice

When the same event study is estimated on the raw acquirer panel, where the measurement boundary shifts at the deal date, both the TFPR and markup effects reverse sign. Section 3.3.1 characterized the consolidation bias that contaminates this comparison and showed that its aggregate direction is an empirical question. We now assess it by estimating equation (9) on the acquirer-only panel, where pre-merger observations reflect the acquirer alone and post-merger observations reflect the merged entity.

Figure 3 reports acquirer-only TFPR event studies. Gross-output TFPR on the acquirer-only panel declines by approximately 2% in the first several post-merger years before gradually recovering toward zero by year eight (Panel A). On the boundary-consistent unit, gross-output TFPR rises steadily over the same horizon, reaching approximately 4% by year eight (Section 4.1). The acquirer-only specification reverses the sign of the estimated effect throughout most of the post-merger period. Value-added TFPR shows a similar pattern, remaining flat to slightly negative on the acquirer-only panel throughout most of the post-merger horizon while rising to approximately 9% on the boundary-consistent unit (Panel B). The gap between the two specifications at any given horizon reflects the aggregate consolidation bias from (8). Under constant returns to scale, the deal-level bias is downward whenever the acquirer is at least as productive as the target (Proposition 1). The aggregate bias in the sample is large enough to reverse the sign of the TFPR estimate. The boundary correction does not merely adjust the magnitude of the TFPR effect. It changes the sign.

Figure 3 also reports acquirer-only markup event studies. The pattern mirrors the TFPR reversal. The inverse materials share, which rises by approximately 6% on the boundary-consistent unit, turns negative on the acquirer-only panel, with point estimates around -1% to -4% in the post-merger period (Panel C). The production-function markup shows a similar downward shift (Panel D). The markup boundary bias at the deal level equals $w_T(\mu_T - \mu_A)$, as established in Section 3.3.1. The direction of the aggregate bias indicates that targets tend to have lower markups than acquirers in the sample, so the boundary shift at the deal date pulls the pooled markup below the acquirer's standalone level. The boundary correction does not merely sharpen the markup estimates. It reverses their direction.

The problem is general to any study that tracks an acquirer through a merger in firm-level accounting data. When one firm absorbs another, the target's operations are folded into the acquirer's financial statements, changing what the acquirer's accounts measure. The comparison of the acquirer before and after the deal is a comparison of two different entities, and the resulting bias is not a second-order adjustment. In this sample, it is large enough to reverse qualitative conclusions about both productivity and markups. Virtual consolidation is a prerequisite for credible inference whenever the unit of observation changes at the deal date.

4.5 The Pre-Merger Aggregation Bias in Practice

Virtual consolidation pools accounts before computing markups, and the resulting combined pre-merger markup is implicitly cost-weighted, equal to the cost-weighted average of the individual markups (Section 3.3.3). If pre-merger markups are instead computed separately for each firm and averaged with sales weights, the pre-merger level is systematically overstated (Proposition 2). When the event study is re-estimated using sales-weighted pre-merger markups on the same boundary-consistent unit, the estimated post-merger markup effect reverses sign (Figure 4). The inverse materials share, which rises by approximately 6% under cost-weighted aggregation (Section 4.2), falls by approximately 5% under sales-weighted aggregation (Panel A). The production-function markup shows the same reversal, falling by approximately 5 to 7% compared

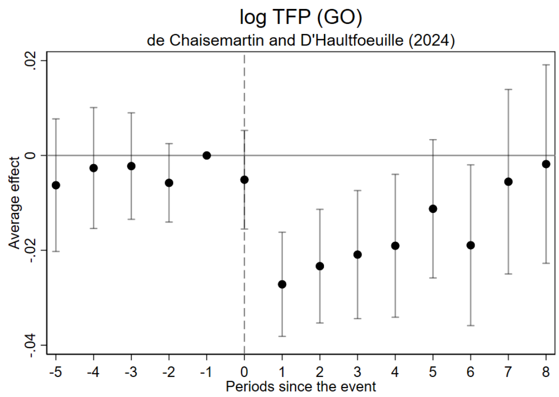

(a) Gross-Output TFPR

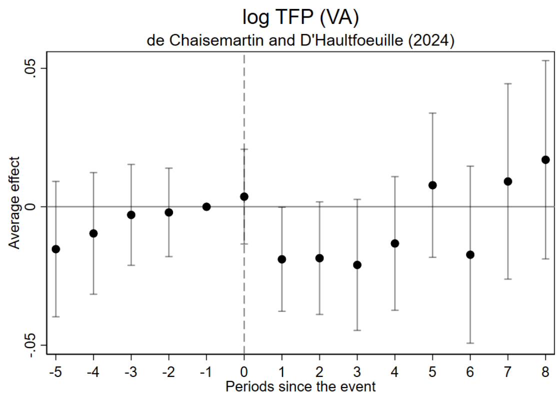

(b) Value-Added TFPR

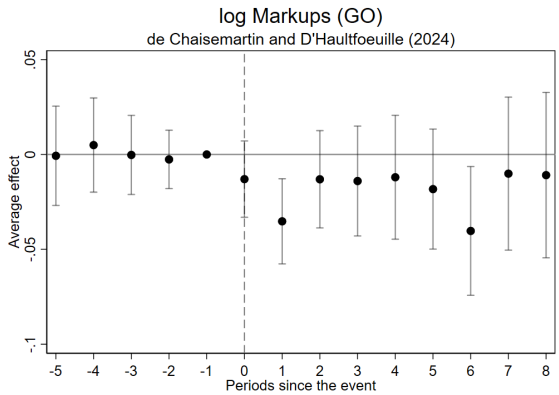

(c) Inverse Materials Share

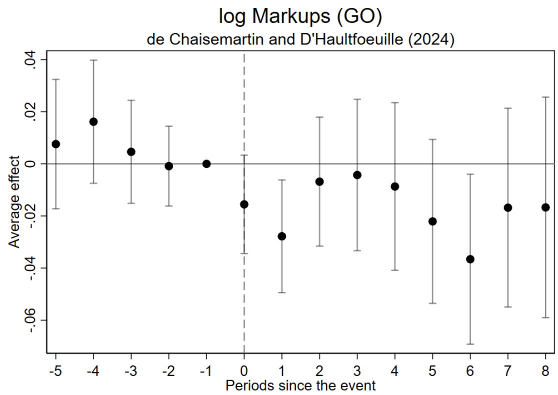

(d) Production-Function Markup

Figure 3: Effect of M&A on TFPR and Markups (Acquirer-Only Panel)

Note: Event-study estimates of the effect of horizontal M&A on log TFPR (Panels A–B) and markups (Panels C–D) for acquirers only, without pre-merger pooling of acquirer and target accounts. Panel C reports the inverse materials share $\ln(Q_{it}/M_{it})$. Panel D reports the production-function markup using fitted gross output. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used. The control group includes firms never involved in a merger and not-yet-treated pairs. Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figures 1 and 2, which report the same specifications on the boundary-consistent acquirer–target unit.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

to a 3% increase under cost-weighted aggregation (Panel B). The aggregation bias alone is large enough to change the qualitative conclusion.

The bias operates through an asymmetry between the pre- and post-merger periods. Before the deal, both firms report separately, and the sales-weighted average of their markups exceeds the markup from pooled accounts by Jensen's inequality (Proposition 2). After the deal, the target ceases to report. There is only one entity, and its markup is simply the merged firm's markup, the same object reported in the baseline results. The sales-weighted specification therefore overstates the pre-merger level but leaves the post-merger level unchanged, understating the post-merger markup increase. The drop at the deal date reflects the disappearance of this wedge, not a decline in markups.

The divergence between the two specifications appears as a level shift at the deal date and remains

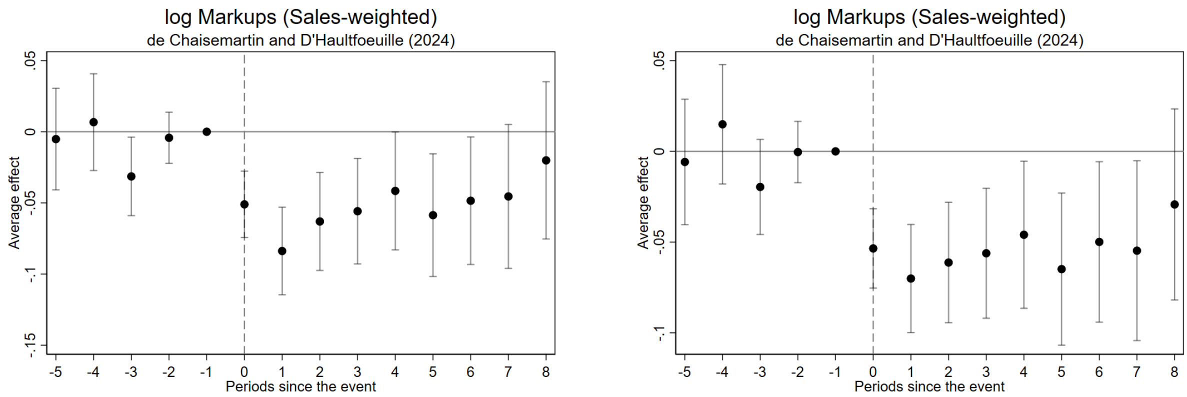

(a) Inverse Materials Share

(b) Production-Function Markup

Figure 4: Effect of M&A on Markups (Sales-Weighted Pre-Merger Aggregation)

Note: Event-study estimates of the effect of horizontal M&A on markups for the boundary-consistent acquirer–target unit, using sales-weighted aggregation of pre-merger markups. Panel A reports the inverse materials share $\ln(Q_{it}/M_{it})$, which under Cobb–Douglas is proportional to the log markup. Panel B reports the production-function markup using fitted gross output. Pre-merger markups are computed as $\mu_t^{\text{sales}} = \pi_{A,t}\,\mu_{A,t} + \pi_{T,t}\,\mu_{T,t}$, where $\pi_{i,t} \equiv Q_{i,t}/(Q_{A,t} + Q_{T,t})$ are sales weights. Post-merger markups are the merged entity's markup. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used. The control group includes firms never involved in a merger and not-yet-treated pairs. Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figure 2, which reports the same specifications using cost-weighted (pooled-accounts) aggregation.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

roughly stable across post-merger horizons. Because both specifications measure the same object after the deal — the merged entity's markup — any post-merger changes affect both equally, and the gap between them is determined entirely by the pre-merger baseline. After the initial shift, both specifications trend upward at similar rates, and the markup dynamics do not differ across the two aggregation methods. The aggregation bias and the boundary bias (Section 4.4) are each large enough on their own to reverse the sign of the estimated markup effect.

5 Robustness

The baseline estimates could reflect selection into treatment, choices in how inputs and outputs are measured, or decisions in sample construction. We examine each concern in turn; the main findings are robust to each. Figures for all robustness checks are reported in Appendix F. Results under alternative estimators and control groups appear in Appendix D.

5.1 Matched Difference-in-Differences

Firms involved in mergers are systematically larger than non-merging firms in total assets, employment, and sales. If these size differences correlate with outcome dynamics, the baseline event-study estimates could reflect compositional differences between treated and control groups rather than merger effects. We address this with a propensity-score-matched difference-in-differences design. Using only pre-treatment observations, we collapse the data to the firm level and estimate propensity scores based on log total assets, tangible fixed assets, labor costs, material costs, and sales. Each treated pair is matched to the nearest control firm (one-to-one nearest-neighbor matching with a caliper of 0.01), and we re-estimate the event studies on the matched sample.

All four outcomes replicate (Figure F.1). Gross-output TFPR rises to approximately 4.5% by year eight, value-added TFPR to approximately 9.5%, and the inverse materials share to approximately 7.5%, all closely tracking the baseline magnitudes and dynamics. The production-function markup follows the same directional pattern at approximately 3–4%. Pre-event estimates are close to zero across all four outcomes. The minor pre-trend departures present in the baseline, at $\ell = -3$ for gross-output TFPR and $\ell = -4$ for the production-function markup, disappear in the matched sample. The results are not driven by observable compositional differences between merging and non-merging firms.

5.2 Four-Digit Industry Classification

The baseline defines horizontal mergers as those in which the acquirer and target share the same NACE 2-digit industry. This classification is broad, and some deals it captures may involve firms in different product markets within the same 2-digit code (for example, NACE 46 includes both wholesale of grain and animal feeds (46.21) and wholesale of computers and software (46.51)). We re-estimate the event studies restricting to mergers in which both firms share the same NACE 4-digit industry, reducing the sample from 3,759 to 2,809 deals. All four outcomes survive (Figure F.2). Gross-output TFPR rises to approximately 4% by year eight, and the inverse materials share to approximately 7%, both essentially unchanged from the baseline. Value-added TFPR reaches approximately 7.5% by year eight, slightly below the baseline terminal estimate of 9%. The production-function markup follows the same directional pattern. The results hold under a substantially narrower definition of within-industry mergers.

5.3 Alternative Output and Capital Measures

The TFPR and markup estimates depend on how output and capital are measured in the production function, and different choices could yield different residuals. The baseline uses gross output, which captures current-period production and is standard for production function estimation with firm-level data. Sales can differ from gross output due to changes in inventories of finished goods. Total assets is used as the capital proxy, which captures acquired intangibles that arise in the M&A process and can have significant measurement implications when omitted (Hess, Rojas, Santaeulalia-Llopis, & Villegas-Sanchez, 2025). We re-estimate using sales as the output measure and, separately, tangible fixed assets as the capital proxy. Under the sales specification, TFPR estimates are smaller, consistent with sales being a noisier output measure, while both markup measures are comparable to the baseline (Figure F.3). Under tangible fixed assets, gross-output TFPR and both markup measures are preserved, though value-added TFPR is smaller (Figure F.4). All outcomes preserve their sign under both alternatives.

5.4 Inclusion of Cross-Industry Mergers

The baseline restricts to horizontal mergers, where acquirer and target share the same NACE 2-digit industry, because the virtual consolidation methodology requires a common production technology to construct the combined unit. This excludes cross-industry deals from the sample. We re-estimate including all M&A deals, assigning the acquirer's industry-level output elasticities to cross-industry pairs, which more than doubles the number of mergers. All four outcomes survive (Figure F.5). TFPR effects are comparable to the baseline. Both markup measures are if anything larger when cross-industry deals are included. The larger markup estimates should be interpreted with caution, however, since the single-technology markup formula may not apply cleanly when acquirer and target operate under different production technologies.

5.5 Excluding Misrecorded Deal Dates

Orbis M&A occasionally records announcement dates with a lag, so some targets disappear from the firm panel before the recorded deal year. The baseline treats these cases as misreporting and reassigns the deal year to the year following the target's last appearance in the data. We implement a more conservative approach by dropping deals where the target disappears two or more years before the recorded announcement rather than correcting the date. All four outcomes survive with comparable magnitudes (Figure F.6).

6 Conclusion

We study how the merged firm's revenue productivity and markups evolve after mergers broadly across industries. We construct the combined acquirer-target entity before each deal by pooling financial statements, maintaining a constant measurement boundary throughout the event window, and estimate event-study regressions across 3,759 horizontal mergers in Spain. We formalize the biases that arise without this correction and show that they are large enough to reverse both estimated effects. On the boundary-consistent unit, revenue productivity rises by approximately 4 percent within eight years of the merger and markups by approximately 5 to 6 percent. These effects build gradually, a pattern more consistent with operational integration than with immediate financial restructuring.

The industry decomposition reveals that these two increases are not alike. Revenue productivity gains are broad-based across industries, while the overall markup increase is driven by large increases in a few industries, particularly human health activities, rather than by moderate increases across many. In the majority of industries where mergers occur, markup effects are small or negative.

These findings put the existing evidence in context. The industry-specific literature has produced conflicting results about whether mergers improve efficiency or increase market power. The representative-economy benchmark suggests that this disagreement reflects genuine variation across industries. Researchers studying a particular industry can now assess whether their findings align with this benchmark or represent an industry-specific departure from it. Beyond the empirical findings, the measurement framework extends to other settings. The consolidation bias arises in any study that tracks an acquirer through a merger in firm-level accounting data, and the virtual consolidation correction requires only that both firms' financial statements be observable before the deal. The approach can be applied in any country with adequate coverage in Orbis or comparable databases.

The principal limitation of this study is that revenue data cannot separate physical efficiency from pricing. The cross-industry variation provides suggestive evidence, but a definitive decomposition requires quantity data at economy-wide scale. Such data are not currently available. Until then, the representative-economy benchmark documented here provides the broadest available picture of what the merged firm becomes.

References

Ackerberg, D., Caves, K., & Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, *83*(6), 2411–2451. [10, 38]

Ackerberg, D., & De Loecker, J. (2024). Production function identification under imperfect competition. Working paper. [10, 38]

Ashenfelter, O. C., Hosken, D. S., & Weinberg, M. C. (2015). Efficiencies brewed: pricing and consolidation in the us beer industry. *The RAND Journal of Economics*, *46*(2), 328–361. [5]

Asker, J., & Nocke, V. (2021). Collusion, mergers, and related antitrust issues. In *Handbook of industrial organization* (Vol. 5, pp. 177–279). Elsevier. [5]

Bhattacharya, V., Illanes, G., & Stillerman, D. (2023). *Merger effects and antitrust enforcement: Evidence from us consumer packaged goods* (NBER Working Paper No. 31123). National Bureau of Economic Research. [5]

Blonigen, B. A., & Pierce, J. R. (2016). *Evidence for the effects of mergers on market power and efficiency* (NBER Working Paper No. 22750). National Bureau of Economic Research. [2, 5, 16, 17]

Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*, rdae007. [15, 16, 48, 49, 50, 51]

Braguinsky, S., Ohyama, A., Okazaki, T., & Syverson, C. (2015). Acquisitions, productivity, and profitability: evidence from the japanese cotton spinning industry. *American Economic Review*, *105*(7), 2086–2119. [5]

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, *225*(2), 200–230. [15, 16, 48, 49, 50, 51]

De Chaisemartin, C., & d'Haultfoeuille, X. (2024). Difference-in-differences estimators of intertemporal treatment effects. *Review of Economics and Statistics*, 1–45. [15, 17, 18, 19, 22, 23, 44, 46, 47, 52, 53, 54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67]

De Chaisemartin, C., & d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American economic review*, *110*(9), 2964–2996. [15, 48]

De Loecker, J., Goldberg, P. K., Khandelwal, A. K., & Pavcnik, N. (2016). Prices, markups, and trade reform. *Econometrica*, *84*(2), 445–510. [10, 39]

De Loecker, J., & Warzynski, F. (2012). Markups and firm-level export status. *American economic review*, *102*(6), 2437–2471. [2, 8]

Demirer, M., & Karaduman, Ö. (2024). *Do mergers and acquisitions improve efficiency? evidence from power plants* (NBER Working Paper No. 32727). National Bureau of Economic Research. [2, 5, 16]

Fons-Rosen, C., Kalemli-Ozcan, Ş., Sørensen, B. E., Villegas-Sanchez, C., & Volosovych, V. (2021). Quantifying productivity gains from foreign investment. *Journal of International Economics*, *131*, 103456. doi: 10.1016/j.jinteco.2021.103456 [16]

Gandhi, A., Navarro, S., & Rivers, D. A. (2020). On the identification of gross output production functions. *Journal of Political Economy*, *128*(8), 2973–3016. [10]

Hess, J. T., Rojas, L., Santaeulalia-Llopis, R., & Villegas-Sanchez, C. (2025). Profits, labor share and the rise of acquired intangibles. Working paper. [24]

Kalemli-Özcan, Ş., Sørensen, B. E., Villegas-Sanchez, C., Volosovych, V., & Yeşiltaş, S. (2024). How to construct nationally representative firm-level data from the orbis global database: New facts on smes and aggregate implications for industry concentration. *American Economic Journal: Macroeconomics*, *16*(2), 353–374. [6, 34, 36]

Kulick, R. B. (2017). *Ready-to-mix: horizontal mergers, prices, and productivity* (Working Paper No. CES-WP-17-38). US Census Bureau Center for Economic Studies. [5]

Miller, N. H., & Weinberg, M. C. (2017). Understanding the price effects of the millercoors joint venture. *Econometrica*, *85*(6), 1763–1791. [2, 5]

Olley, G. S., & Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica: Journal of the Econometric Society*, 1263–1297. [39]

Stiebale, J., & Szücs, F. (2022). Mergers and market power: evidence from rivals' responses in european markets. *The RAND Journal of Economics*, *53*(4), 678–702. [5, 17]

Stiebale, J., & Vencappa, D. (2018). Acquisitions, markups, efficiency, and product quality: Evidence from india. *Journal of International Economics*, *112*, 70–87. [2, 5, 16, 17]

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, *225*(2), 175–199. [15, 16, 48, 49, 50, 51]

Syverson, C. (2019). Macroeconomics and market power: Context, implications, and open questions. *Journal of Economic Perspectives*, *33*(3), 23–43. [18]

Measuring Merger Effects with Revenue Data

Jacob T. Hess and Xufeng Wang

Online Appendix (Not For Publication)

Appendix A Data

A.1 Sample Description

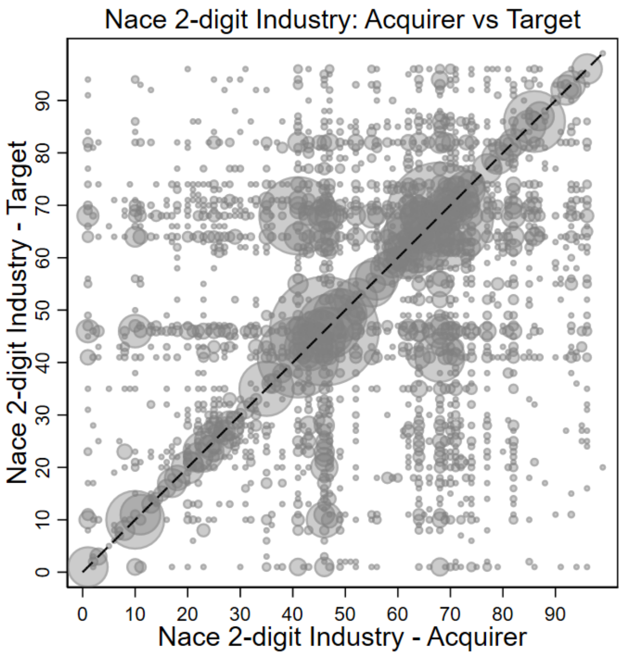

Figure A.1: Acquirer and Target Industries (NACE 2-Digit)

Note: Each bubble represents an acquirer–target industry pair, with bubble size proportional to the number of transactions. The dashed diagonal indicates within-industry (horizontal) deals. More than half of all transactions fall off the diagonal, involving acquirers and targets in different 2-digit industries.

Source: Authors' calculations using Orbis Spain and Orbis M&A.

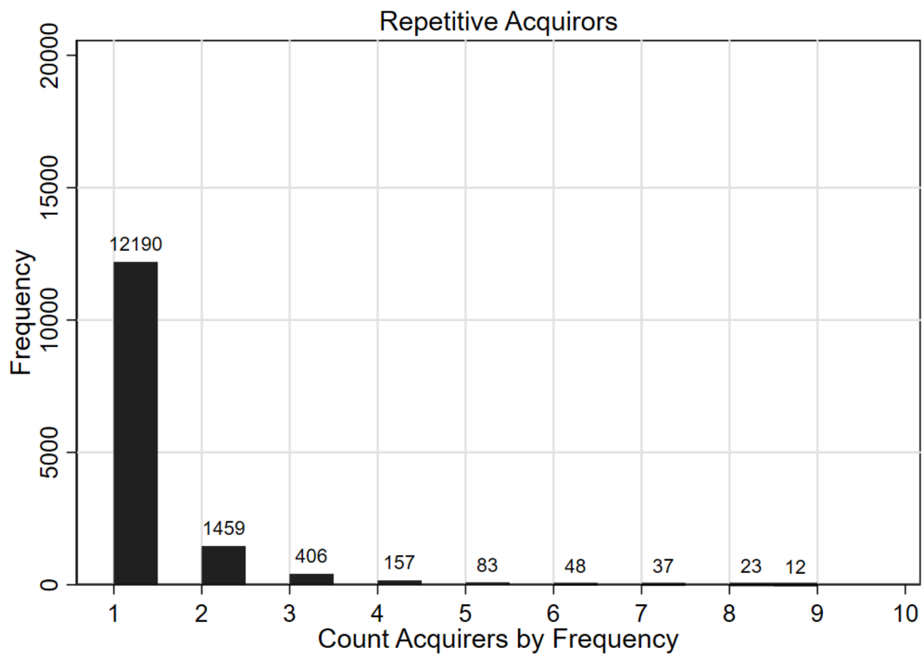

Figure A.2: Distribution of Acquirers by Number of Transactions

Note: The figure shows the number of distinct acquirers by their total count of M&A transactions during 1997–2022. Approximately 85% of acquirers participate in a single transaction. The baseline analysis restricts attention to single-transaction firms.

Source: Authors' calculations using Orbis Spain and Orbis M&A.

A.2 Representativeness of Orbis Spain

Orbis Spain covers approximately 50% of aggregate value added and 52% of labor compensation over 1997–2022, compared to Eurostat national accounts (Table A.1). Coverage improves substantially when public-service sectors (NACE Rev. 2 codes above 82) are excluded, rising to 60% for value added and 72% for labor compensation. Employment coverage is 41% of total employment and 48% excluding the self-employed, who account for about one-sixth of total employment and are not captured in Orbis (Table A.2). Excluding both public services and the self-employed, employment coverage reaches 65%.

The sample covers approximately 64% of all registered firms in Spain, averaged over 1999–2022 (Table A.3).[8] The firm size distribution in Orbis closely matches the national distribution: 83% of firms in Orbis have fewer than 10 employees, compared to 88% in the INE data, and the shares are similar across all size classes.

Tables A.4–A.6 report coverage by one-digit NACE sector. Private-sector coverage is generally strong, particularly in industry, information and communication, and professional services. Coverage of public-service sectors (public administration, education, and health) is limited, consistent with the aggregate patterns above. Some sectors show coverage exceeding 100%,

[8] We use INE instead of Eurostat for the firm count comparison because the latter only provides aggregate number of firms starting in 2005.

because firms operating across multiple sectors are assigned a single core NACE code in Orbis, while Eurostat may allocate activity across sectors.

Table A.1: Orbis Spain Coverage of Value Added and Labor Compensation (1997–2022)

	Orbis	Eurostat	Share (%)
Value Added (bil.)	464.2	938.5	49.5
Value Added (bil.) - NACE Rev. 2 $\leq$ 82	442.9	732.6	60.4
Labor Compensation (bil.)	254.2	489.1	52.0
Labor Compensation (bil.) - NACE Rev. 2 $\leq$ 82	237.8	330.4	72.0

Note: The table compares averages (in billions of euros) from Orbis Spain and Eurostat over the years 1997 to 2022. All variables are in constant 2015 prices. Value added and labor compensation from Orbis are the unweighted sums by year and averaged over the time period, whereas the aggregate variables are taken from the national accounts in Eurostat (table code: nama_10_a10).
Source: Authors' calculations using Orbis Spain and Eurostat.

Table A.2: Orbis Spain Coverage of Employment (1997–2022)

	Total			No Self-Employed	
	Orbis	Eurostat	Share (%)	Eurostat	Share (%)
Emp. (mil.)	7.6	18.7	40.6	15.9	47.8
Emp. (mil.) NACE Rev. 2 $\leq$ 82	7.1	13.3	52.9	10.9	65.0

Note: The table compares total number of employees (in millions) from Orbis Spain and Eurostat (table code: nama_10_a10_e) averaged over the years 1997 to 2022. The second, third, and fourth columns compare Orbis employment to the total number of individuals employed. Orbis by construction does not contain self-employed individuals. Columns 5 and 6 show total employment when the self-employed are excluded.
Source: Authors' calculations using Orbis Spain and Eurostat.

Table A.3: Orbis Spain Coverage of the Firm Size Distribution (1999–2022)

	Orbis		INE	
Employees	Nbr	Share (%)	Nbr	Share (%)
0 to 9	590,195	83.0	978,854	87.8
10 to 19	63,385	8.9	72,801	6.5
20 to 49	39,232	5.5	41,899	3.8
50 to 199	14,778	2.1	16,634	1.5
200+	3,878	0.5	4,365	0.4
Total	711,469	100.0	1,114,553	100.0

Note: The table compares the firm size distribution by employment in Orbis Spain with the aggregate number of firms from the Central Business Register in INE. Firms consist of those either publicly listed or with limited liability. The distributions are averaged over the years 1999–2022. Self-employed are not included.
Source: Authors' calculations using Orbis Spain and National Statistics Institute.

A.2.1 Sector-Level Coverage

Table A.4: Orbis Spain Coverage of Value Added by Sector (1997–2022)

	Orbis	Eurostat	Share (%)
(A) Agriculture	5.1	29.5	17.2
(B–E) Industry (except construction)	133.7	164.9	81.1
(F) Construction	42.8	78.6	54.5
(G–I) Trade, transport, and accommodation	118.3	215.7	54.8
(J) Information and communication	32.5	37.9	85.6
(K) Financial and insurance activities	44.2	40.8	108.2
(L) Real estate activities	15.5	93.1	16.6
(M–N) Professional and administrative services	50.8	72.1	70.5
(O–Q) Public administration, education, and health	12.0	164.5	7.3
(R–U) Arts, entertainment, and other services	8.6	41.4	20.7
Total	463.4	938.5	49.4

Note: The table compares value added (in billions of euros) by NACE Rev. 2 section from Orbis Spain and Eurostat averaged over the years 1997 to 2022. All variables are in constant 2015 prices. Value added from Orbis is the unweighted sum by year and averaged over the time period, whereas the aggregate variables are taken from the national accounts in Eurostat (table code: nama_10_a10). Sector names follow NACE Rev. 2 sections, abbreviated for readability.
Source: Authors' calculations using Orbis Spain and Eurostat.

Table A.5: Orbis Spain Coverage of Labor Compensation by Sector (1997–2022)

	Orbis	Eurostat	Share (%)
(A) Agriculture	3.0	5.0	59.4
(B–E) Industry (except construction)	68.7	83.9	81.8
(F) Construction	27.4	40.8	67.2
(G–I) Trade, transport, and accommodation	78.3	111.0	70.5
(J) Information and communication	17.4	20.1	86.5
(K) Financial and insurance activities	4.0	20.7	19.2
(L) Real estate activities	4.7	3.5	133.7
(M–N) Professional and administrative services	34.4	45.3	75.9
(O–Q) Public administration, education, and health	9.9	132.9	7.5
(R–U) Arts, entertainment, and other services	5.2	25.9	20.1
Total	253.0	489.1	51.7

Note: The table compares labor compensation (in billions of euros) by NACE Rev. 2 section from Orbis Spain and Eurostat averaged over the years 1997 to 2022. All variables are in constant 2015 prices. Labor compensation from Orbis is the unweighted sum by year and averaged over the time period, whereas the aggregate variables are taken from the national accounts in Eurostat (table code: nama_10_a10). Sector names follow NACE Rev. 2 sections, abbreviated for readability.
Source: Authors' calculations using Orbis Spain and Eurostat.

Table A.6: Orbis Spain Coverage of Employment by Sector (1997–2022)

	Total			No Self-Employed	
	Orbis	Eurostat	Share (%)	Eurostat	Share (%)
(A) Agriculture	0.1	0.8	16.1	0.5	30.3
(B–E) Industry (except construction)	1.7	2.6	67.7	2.4	71.5
(F) Construction	0.8	1.7	49.0	1.4	57.4
(G–I) Trade, transport, and accommodation	2.5	5.3	47.4	4.1	62.4
(J) Information and communication	0.3	0.4	79.7	0.4	86.6
(K) Financial and insurance activities	0.1	0.4	39.3	0.3	42.2
(L) Real estate activities	0.1	0.2	81.5	0.1	123.7
(M–N) Professional and administrative services	1.2	1.9	62.0	1.7	73.0
(O–Q) Public administration, education, and health	0.4	3.8	9.3	3.7	9.5
(R–U) Arts, entertainment, and other services	0.2	1.6	10.8	1.3	12.8
Total	7.6	18.7	40.5	15.9	47.7

Note: The table compares total number of employees (in millions) by NACE Rev. 2 section from Orbis Spain and Eurostat (table code: nama_10_a10_e) averaged over the years 1997 to 2022. The second, third, and fourth columns compare Orbis employment to the total number of individuals employed. Orbis by construction does not contain self-employed individuals. Columns 5 and 6 show total employment when the self-employed are excluded. Sector names follow NACE Rev. 2 sections, abbreviated for readability.
Source: Authors' calculations using Orbis Spain and Eurostat.

A.3 Sample Construction

This subsection documents the construction of the analysis sample from the raw Orbis downloads to the final event-study panel. The steps proceed from the firm-level panel to the M&A transactions, variable cleaning, and the construction of acquirer-target pairs.

A.3.1 Orbis Spain Panel

We downloaded Orbis Historical from Moody's DataHub in August 2024. The following steps produce the baseline firm-year panel.

1. We construct the operating year following Kalemli-Özcan et al. (2024). If the closing date is on or after June 1st, the current year is assigned as the operating year. If the closing date is before June 1st, the previous year is assigned. The BvD ID number and year serve as unique identifiers in the Orbis dataset.

2. We drop observations with the consolidated code "C2", which denotes a statement of a legal entity integrating the statements of its controlled subsidiaries or branches. These firms have two accounts: one for the entire enterprise (headquarters and all branches) and one for the headquarters. We keep the unconsolidated headquarters account (U2) when both exist.

3. We keep only observations where the reporting period equals 12 months to ensure comparability across firm-years.

4. After completing the above steps, the dataset still contains a small portion (0.029%) of duplicate observations on the BvD ID number and year identifiers. These duplicates arise from reporting on different closing dates or through different filing types (local registry filings or annual reports). In these cases, we keep the record with the latest date for the reporting year.

5. We drop the years before 1997 because M&A data starts from 1997, as well as the year 2023, which has only 10 observations after the above cleaning steps.

All monetary variables are deflated to constant 2015 prices using the GDP deflator from the World Bank's World Development Indicators.[9]

[9] https://databank.worldbank.org/source/world-development-indicators

Table A.7: Sample Attrition (Orbis Spain Panel)

Cause	# Obs Lost	Proportion of Total %	Remaining
# ObS in Raw Data			19,131,872
Step 2: Drop consolidated code C2	62,668	0.325	19,069,204
Step 2: Drop consolidated code C1	5,008	0.026	19,064,196
Step 3: Drop number_of_months $\neq$ 12	712,195	3.723	18,352,001
Step 4: Drop duplicates - keeping latest	2,063	0.011	18,349,938
Step 5: Drop years	698,974	3.665	17,650,964
N. Remaining in Sample	1,480,908	7.741	17,650,964

A.3.2 M&A Transactions

We obtained M&A transaction records from Orbis M&A (formerly Zephyr) in August 2024. The raw data for Spain contains 40,636 deals, of which 40,180 are classified as acquisitions and 456 as mergers. Since mergers account for only about 1% of all deals, we do not distinguish between the two and refer to all transactions collectively as M&As.

Orbis M&A reports different dates depending on when the deal status is verified. We use the announcement date as the baseline event year. However, for approximately 50% of deals the announcement date coincides with the completion date, likely due to delays in recording the announcement date. We shift the announcement date forward by one year when it is identical to the completion date.

We link deals to the firm panel using BvD ID numbers and drop observations with missing identifiers. The dataset includes deals with various statuses (“rumored,” “pending,” “withdrawn”); we retain only deals with a status of “completed” or “completed assumed.” We restrict attention to domestic transactions in which both the acquirer and the target are Spanish firms, and to deals with a single acquirer and a single target. Firms in the panel that are involved in non-qualifying transactions (foreign parties, multiple acquirers or targets) are also dropped from the analysis sample.

We then use information from the Orbis Spain panel to correct the event year. Table A.8 reports the time difference between the last year a target firm appears in Orbis Spain and the M&A year recorded in Orbis M&A. For cases in which the target disappeared one, two, or three years before the recorded M&A year, we redefine the event year as the year immediately following the target’s last appearance in the data.

After applying these restrictions, the sample contains 8,655 domestic M&A transactions covering 1997–2022. Of these, 3,759 are horizontal deals in which the acquirer and target operate within the same NACE 2-digit industry; these form the baseline event-study sample. The remaining cross-industry deals are excluded from the baseline but included in a robustness check (Section 5).

Table A.8: Target Disappearance Relative to M&A Announcement Date

Relative Year	Observations
−5	158
−4	269
−3	494
−2	2,109
−1	5,997
0	3,773
1	530
2	273
3	246
4	216
5	179
6	143
7	120
8	92
Total	15,613

Note: The relative year is the difference between the last year the target firm appears in Orbis Spain and the M&A announcement year recorded in Orbis M&A. A value of −1 means the target's last observation was one year before the recorded announcement. Cases where the target's last year in Orbis Spain is 2022 are excluded to avoid confounding disappearance with the end of the sample period.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

A.3.3 Variable Cleaning

Following the guidelines in Appendix 5.3 of Kalemli-Özcan et al. (2024), we apply the following variable-level cleaning steps to the merged dataset.

1. We remove company-years that lack simultaneous data for total assets, operating revenue, sales, and employment.

2. We exclude companies entirely if they report negative values for total assets, materials, fixed assets, or depreciation and amortization in any year.

3. We drop firms if employment or wage bills are negative in any year.

4. We remove firms reporting negative costs of goods sold, sales, or tangible fixed assets in any year.

5. We drop all remaining observations in which any of the following variables are non-positive or missing: value added, sales, fixed assets, total assets, number of employees, labor costs, and material costs.

6. We exclude firms if their employment per million of total assets falls below 0.1 or exceeds the 99.9th percentile. Firms with employment per million of sales or sales-to-total-assets ratios exceeding the 99.9th percentile are also removed.

After these steps, the number of observations decreases from 17,206,941 to 7,865,942.

A.3.4 Acquirer-Target Pairs

For each horizontal transaction involving a single-transaction acquirer and target, we construct the paired unit by aggregating the financial statements of both firms in each year of the panel. When data for both the acquirer and the target are available before and after the deal, the paired unit reflects both firms throughout. In cases where only acquirer data are available after the deal, we use the acquirer's accounts as the post-deal observation, on the assumption that the target has been integrated into the acquirer's reporting entity. We apply the same variable cleaning steps described in Appendix A.3.3 to the paired data before merging it with the main panel.

A.3.5 Event-Study Sample

Table A.9 summarizes the final event-study sample.

Table A.9: Event-Study Sample

	Firm-Year Observations	Distinct Firms
Non-M&A firms	7,751,974	1,044,359
Acquirer-target pairs (horizontal M&A)	40,008	3,759

Note: The table reports the number of observations in the analysis sample after merging the Orbis Spain panel with the virtually consolidated acquirer-target pairs. Non-M&A firms are those not involved in any transaction during the sample period. Distinct firms reports the number of unique firms (or pairs) within each group.

Source: Authors' calculations using Orbis Spain and Orbis M&A.

Appendix B Production Function Output Elasticity Estimates

B.1 Estimation procedure

We estimate output elasticities using the control-function approach of Ackerberg et al. (2015). We present the procedure for the gross-output specification used in the main text. The value-added specification follows the same steps with output y_{it} replacing q_{it} and the materials input dropped.

First stage. The gross-output production function in logs is

$$\begin{aligned} q_{it} &= \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \omega_{it} + \epsilon_{it} \\ &= \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + g_t(k_{it}, m_{it}, l_{it}, E_{it}) + \epsilon_{it} \\ &= \Phi(k_{it}, l_{it}, m_{it}, E_{it}) + \epsilon_{it}, \end{aligned}$$

where the second equality uses the assumption that materials demand is strictly increasing in ω_{it} conditional on capital, labor, and observed demand shifters (firm location, firm type, listing status, and year fixed effects), so that $\omega_{it} = g_t(k_{it}, m_{it}, l_{it}, E_{it})$. We approximate $\Phi(k_{it}, l_{it}, m_{it}, E_{it}) \equiv \beta_0 + \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \omega_{it}$ with a third-order polynomial in (k_{it}, l_{it}, m_{it}) augmented with E_{it} and recover $\widehat{\Phi}_{it}$ from OLS.

Second stage. Productivity follows a first-order Markov process, $\omega_{it} = \rho_1 \omega_{it-1} + \xi_{it}$, where ξ_{it} is the productivity innovation. Substituting the expression for ω_{it} from the first stage yields

$$q_{it} - \beta_0 - \beta_k k_{it} - \beta_l l_{it} - \beta_m m_{it} - \epsilon_{it} = \rho_1 \big(\Phi(k_{it-1}, l_{it-1}, m_{it-1}) - \beta_0 - \beta_k k_{it-1} - \beta_l l_{it-1} - \beta_m m_{it-1} \big) + \xi_{it}.$$

The composite residual $\xi_{it} + \epsilon_{it}$ depends on five parameters (β_0, β_k, β_l, β_m, ρ_1). Following the appendix of Ackerberg et al. (2015), the intercept and the autoregressive coefficient can be concentrated out, reducing the parameter vector to $(\beta_k, \beta_l, \beta_m)$. The output elasticities are identified from the moment conditions

$$E\left[\widehat{\xi}_{it}(\beta_k, \beta_l, \beta_m) \otimes \begin{pmatrix} k_{it} \\ l_{it-1} \\ m_{it-1} \end{pmatrix} \right] = 0,$$

where $\widehat{\xi}_{it}(\beta_k, \beta_l, \beta_m)$ is the residual from regressing $q_{it} - \beta_0 - \beta_k k_{it} - \beta_l l_{it} - \beta_m m_{it}$ on $\Phi(k_{it-1}, l_{it-1}, m_{it-1}) - \beta_0 - \beta_k k_{it-1} - \beta_l l_{it-1} - \beta_m m_{it-1}$ and a constant. Capital k_{it} is predetermined and therefore orthogonal to ξ_{it}. Labor l_{it-1} is a dynamic input subject to adjustment frictions, and its lag is orthogonal to the current innovation. The instrument m_{it-1} is constructed following Ackerberg and De Loecker (2024).

B.2 Selection correction

The baseline estimation excludes all firms involved in M&A. Excluding these firms may introduce selection bias if M&A participation is correlated with unobserved productivity. As a robustness check, we apply the approach of De Loecker et al. (2016), who adapt the Olley and Pakes (1996) selection correction to address selection between single-product and multi-product firms. We use the same procedure to correct for selection between M&A-involved and non-M&A-involved firms.

Firms involved in M&A are plausibly more productive than other firms, conditional on their capital and labor. Our setting differs from standard entry and exit selection, where the participation decision depends on productivity in the current period. For M&A, we assume that the participation decision is based on productivity in the previous period, reflecting the time required to identify, evaluate, and complete a transaction. The productivity threshold at time t is determined by the firm's state variables at $t-1$, as in De Loecker et al. (2016).

Let $\chi_{it+1} = 1$ if firm i remains non-M&A-involved in period $t+1$, and $\chi_{it+1} = 0$ otherwise. The probability of remaining non-M&A-involved satisfies

$$\begin{aligned}\Pr\{\chi_{it+1} = 1 \mid \overline{\omega}_{it+1}, J_t\} &= \Pr\{\omega_{it+1} \leq \overline{\omega}_{it+1} \mid \overline{\omega}_{it+1}, \omega_{it}\} \\ &= \varphi_t(\overline{\omega}_{it+1}, \omega_{it}) \\ &= \varphi_t(m_{it}, k_{it}, l_{it}, E_{it}) \\ &\equiv P_{it},\end{aligned}$$

where J_t is the information set at time t, $\overline{\omega}_{it+1}$ is the productivity threshold, and E_{it} denotes observed demand shifters (firm location, firm type, listing status). The third equality follows from the invertibility of the materials demand function, $\omega_{it} = g_t(k_{it}, m_{it}, l_{it}, E_{it})$.

Since M&A-involved firms account for a small proportion of all firms, we treat the exclusion of these firms as the primary specification and the selection correction as a robustness check. Both approaches yield similar results.

B.3 Estimated output elasticities

Table B.1: Estimated Output Elasticity of Capital and Labor Results

		Value Added		Gross Output		
NACE-2	NACE Description	β_k	β_l	β_k	β_l	β_m
1	Crop and animal production, hunting and related service activities	0.275	0.716	0.118	0.289	0.569
2	Forestry and logging	0.241	0.761	0.123	0.361	0.494
3	Fishing and aquaculture	0.204	0.839	0.132	0.453	0.382
5	Mining of coal and lignite	0.320	0.764	0.296	0.670	0.040
7	Mining of metal ores	0.539	0.449	0.601	0.195	0.087
8	Other mining and quarrying	0.340	0.707	0.386	0.591	0.039
9	Mining support service activities	0.145	0.944	0.104	0.344	0.578
10	Manufacture of food products	0.347	0.486	0.056	0.292	0.666

(Continued on the next page)

(Table continues on the next page)

NACE-2	NACE Description	Value Added β_k	Value Added β_l	Gross Output β_k	Gross Output β_l	Gross Output β_m
11	Manufacture of beverages	0.358	0.686	0.118	0.227	0.673
12	Manufacture of tobacco products	0.664	0.482	0.266	0.692	0.199
13	Manufacture of textiles	0.208	0.827	0.111	0.458	0.445
14	Manufacture of wearing apparel	0.164	0.895	0.424	0.067	0.184
15	Manufacture of leather and related products	0.180	0.849	0.114	0.405	0.480
16	Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials	0.180	0.880	0.074	0.385	0.566
17	Manufacture of paper and paper products	0.330	0.596	0.093	0.392	0.529
18	Printing and reproduction of recorded media	0.211	0.829	0.112	0.437	0.462
20	Manufacture of chemicals and chemical products	0.285	0.765	0.110	0.327	0.588
21	Manufacture of basic pharmaceutical products and pharmaceutical preparations	0.240	0.891	0.445	0.537	0.067
22	Manufacture of rubber and plastic products	0.254	0.785	0.101	0.346	0.553
23	Manufacture of other non-metallic mineral products	0.210	0.843	0.080	0.334	0.596
24	Manufacture of basic metals	0.256	0.776	0.388	0.587	0.200
25	Manufacture of fabricated metal products, except machinery and equipment	0.187	0.855	0.109	0.440	0.471
26	Manufacture of computer, electronic and optical products	0.219	0.802	0.118	0.425	0.458
27	Manufacture of electrical equipment	0.240	0.802	0.139	0.393	0.489
28	Manufacture of machinery and equipment n.e.c.	0.210	0.804	0.172	0.424	0.407
29	Manufacture of motor vehicles, trailers and semi-trailers	0.214	0.806	0.267	0.229	0.388
30	Manufacture of other transport equipment	0.190	0.838			
31	Manufacture of furniture	0.128	0.932	0.036	0.377	0.594
32	Other manufacturing	0.155	0.883	0.050	0.428	0.523
33	Repair and installation of machinery and equipment	0.173	0.858	0.147	0.491	0.365
35	Electricity, gas, steam and air conditioning supply	0.456	0.577	0.406	0.215	0.366
36	Water collection, treatment and supply	0.305	0.725			
37	Sewerage	0.230	0.784	0.250	0.430	0.296

(Continued on the next page)

(Table continues on the next page)

NACE-2	NACE Description	Value Added		Gross Output		
		β_k	β_l	β_k	β_l	β_m
38	Waste collection, treatment and disposal activities; materials recovery	0.266	0.768	0.083	0.457	0.479
39	Remediation activities and other waste management services	0.299	0.714	0.181	0.350	0.435
41	Construction of buildings	0.228	0.781	0.198	0.384	0.412
42	Civil engineering	0.181	0.827	0.166	0.459	0.370
43	Specialised construction activities	0.158	0.878	0.118	0.477	0.412
45	Wholesale and retail trade and repair of motor vehicles and motorcycles	0.156	0.901	0.085	0.280	0.657
46	Wholesale trade, except of motor vehicles and motorcycles	0.253	0.781	0.114	0.260	0.617
47	Retail trade, except of motor vehicles and motorcycles	0.153	0.921	0.057	0.282	0.677
49	Land transport and transport via pipelines	0.228	0.801	0.160	0.430	0.380
50	Water transport	0.244	0.822	0.152	0.438	0.433
51	Air transport	0.350	0.655	0.499	0.465	0.011
52	Warehousing and support activities for transportation	0.267	0.778	0.247	0.384	0.373
53	Postal and courier activities	0.200	0.856	0.165	0.452	0.378
55	Accommodation	0.187	0.816	0.137	0.589	0.252
56	Food and beverage service activities	0.107	0.936	0.057	0.388	0.573
58	Publishing activities	0.230	0.788	0.160	0.459	0.376
59	Motion picture, video and television programme production, sound recording and music publishing activities	0.311	0.713	0.236	0.453	0.318
60	Programming and broadcasting activities	0.321	0.471	0.166	0.584	0.233
61	Telecommunications	0.227	0.809	0.140	0.462	0.413
62	Computer programming, consultancy and related activities	0.193	0.820	0.180	0.543	0.279
63	Information service activities	0.226	0.788	0.220	0.582	0.214
64	Financial service activities, except insurance and pension funding	0.301	0.698	0.255	0.476	0.259
65	Insurance, reinsurance and pension funding, except compulsory social security			0.203	0.611	0.243
66	Activities auxiliary to financial services and insurance activities	0.229	0.784	0.198	0.573	0.232

(Continued on the next page)

(Table continues on the next page)

NACE-2	NACE Description	Value Added		Gross Output		
		β_k	β_l	β_k	β_l	β_m
68	Real estate activities	0.313	0.705	0.265	0.453	0.285
69	Legal and accounting activities	0.196	0.817	0.172	0.643	0.198
70	Activities of head offices; management consultancy activities	0.250	0.764	0.222	0.538	0.247
71	Architectural and engineering activities; technical testing and analysis	0.249	0.752	0.216	0.477	0.292
72	Scientific research and development	0.354	0.496	0.178	0.475	0.319
73	Advertising and market research	0.245	0.779	0.220	0.456	0.350
74	Other professional, scientific and technical activities	0.236	0.793	0.212	0.521	0.283
75	Veterinary activities	0.157	0.873	0.111	0.440	0.444
77	Rental and leasing activities	0.384	0.659	0.304	0.367	0.322
78	Employment activities	0.167	0.853	0.049	0.768	0.282
79	Travel agency, tour operator and other reservation service and related activities	0.217	0.819	0.153	0.346	0.487
80	Security and investigation activities					
81	Services to buildings and landscape activities	0.120	0.902	0.100	0.706	0.206
82	Office administrative, office support and other business support activities	0.212	0.810	0.194	0.576	0.255
84	Public administration and defence; compulsory social security	0.234	0.822	0.216	0.618	0.192
85	Education	0.138	0.895	0.112	0.662	0.227
86	Human health activities	0.254	0.741	0.188	0.492	0.290
87	Residential care activities	0.089	0.939	0.084	0.846	0.093
88	Social work activities without accommodation					
90	Creative, arts and entertainment activities	0.262	0.741	0.204	0.421	0.355
91	Libraries, archives, museums and other cultural activities	0.194	0.854	0.140	0.654	0.228
92	Gambling and betting activities	0.345	0.670	0.263	0.434	0.305
93	Sports activities and amusement and recreation activities	0.178	0.850	0.114	0.624	0.286
94	Activities of membership organisations	0.203	0.797	0.142	0.551	0.309
95	Repair of computers and personal and household goods	0.150	0.911	0.101	0.515	0.410
96	Other personal service activities	0.351	0.419	0.112	0.614	0.334

(Continued on the next page)

(Table continues on the next page)

NACE-2	NACE Description	Value Added β_k	Value Added β_l	Gross Output β_k	Gross Output β_l	Gross Output β_m
97	Activities of households as employers of domestic personnel	0.129	0.936	0.100	0.750	0.214
99	Activities of extraterritorial organisations and bodies	0.047	0.960	0.025	0.807	0.211

Note: The Value Added columns report estimation results based on the value added production function, while the Gross Output columns report results based on the gross output production function. In all specifications, total assets are used as the measure of capital. Output elasticities are set to missing when the estimated elasticity for a given industry is negative.
Source: Author's estimations using Orbis Spain and Orbis M&A.

Appendix C Regression Tables

Table C.1: Event-Study Estimates for Sections 4.1 and 4.2

	TFPR		Markups	
	Gross Output (1)	Value Added (2)	Inverse Materials Share (3)	Production-Function Markup (4)
Time to Treat −5	−0.00179	−0.00223	0.00557	0.0134
	(0.00676)	(0.0141)	(0.0106)	(0.00974)
Time to Treat −4	0.000890	0.00324	0.00983	0.0149*
	(0.00630)	(0.0123)	(0.0101)	(0.00906)
Time to Treat −3	−0.0139**	−0.0118	−0.00583	0.00479
	(0.00545)	(0.0110)	(0.00732)	(0.00644)
Time to Treat −2	−0.00631	−0.00304	−0.000220	0.00458
	(0.00469)	(0.00954)	(0.00670)	(0.00541)
Time to Treat 0	0.0205***	0.0191*	0.0259***	0.00207
	(0.00536)	(0.0104)	(0.00919)	(0.00838)
Time to Treat 1	0.00280	0.0226**	0.0214**	0.00224
	(0.00604)	(0.0110)	(0.0107)	(0.00954)
Time to Treat 2	0.0123*	0.0330***	0.0432***	0.0180
	(0.00632)	(0.0120)	(0.0120)	(0.0113)
Time to Treat 3	0.0227***	0.0453***	0.0477***	0.0186
	(0.00765)	(0.0131)	(0.0140)	(0.0132)
Time to Treat 4	0.0216***	0.0432***	0.0550***	0.0310**
	(0.00832)	(0.0129)	(0.0157)	(0.0156)
Time to Treat 5	0.0235***	0.0611***	0.0540***	0.0247
	(0.00820)	(0.0152)	(0.0156)	(0.0152)
Time to Treat 6	0.0180*	0.0474***	0.0446**	0.0178
	(0.00917)	(0.0182)	(0.0173)	(0.0165)
Time to Treat 7	0.0338***	0.0883***	0.0554***	0.0210
	(0.0110)	(0.0206)	(0.0187)	(0.0174)
Time to Treat 8	0.0437***	0.0893***	0.0685***	0.0290
	(0.0128)	(0.0222)	(0.0244)	(0.0233)
ATT	0.0192***	0.0415***	0.0415***	0.0153
	(0.00489)	(0.00852)	(0.00973)	(0.00942)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Event-time estimates from the De Chaisemartin and d'Haultfoeuille (2024) estimator on the pooled acquirer–target unit. The control group includes firms never involved in a merger and not-yet-treated pairs. Columns (1)–(2) report TFPR from gross-output and value-added production functions (Figure 1). Columns (3)–(4) report the inverse materials share and the production-function markup (Figure 2). Coefficients are log deviations relative to the year before the announcement ($\ell = -1$). Standard errors are clustered at the acquirer–target pair level.

Source: Authors' calculations using Orbis Spain and Orbis M&A.

Appendix D Alternative Specifications

D.1 Alternative Control Group

The main-text results use a combined control group that includes both firms never involved in a merger and not-yet-treated pairs whose deal occurs later in the sample. If never-involved firms differ systematically from eventually merging firms in ways that correlate with outcome dynamics, this could bias the event-study estimates. We address this concern by re-estimating the baseline specification using only not-yet-treated pairs as controls. Because every control unit in this specification is itself a future acquirer–target pair, the comparison is drawn entirely from the population of merging firms.

Figure D.1 reports the results. The TFPR and markup patterns are qualitatively unchanged. Gross-output TFPR rises gradually over the post-merger horizon (Panel A), and value-added TFPR shows a similar trajectory with larger magnitudes (Panel B). The inverse materials share increases over the post-merger period (Panel C), as does the production-function markup (Panel D). The choice of control group does not drive the main findings.

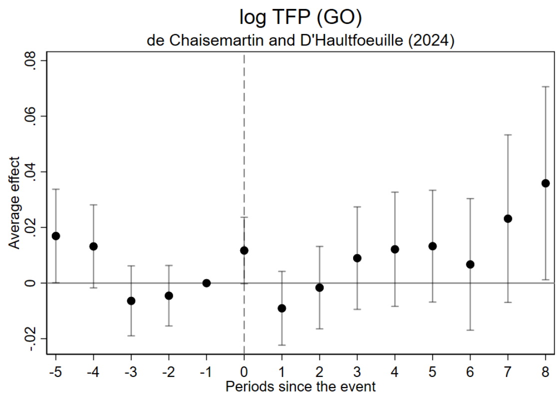

(a) Gross-Output TFPR

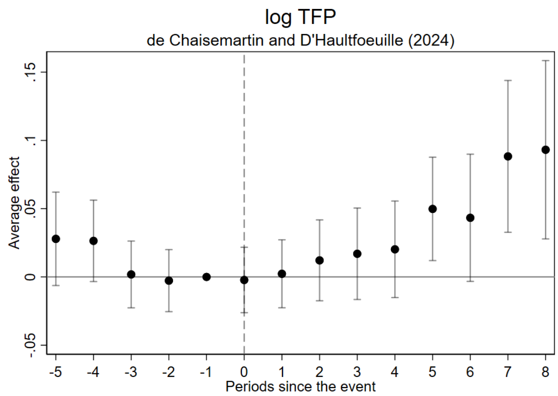

(b) Value-Added TFPR

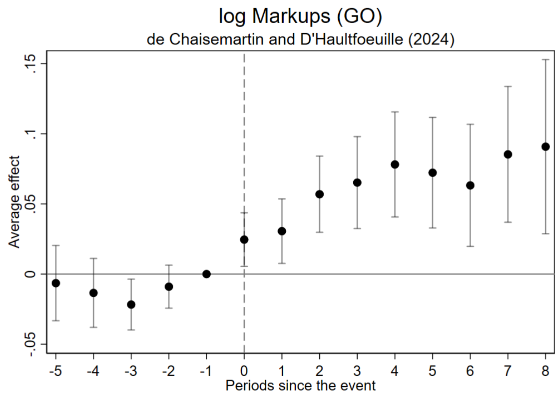

(c) Inverse Materials Share

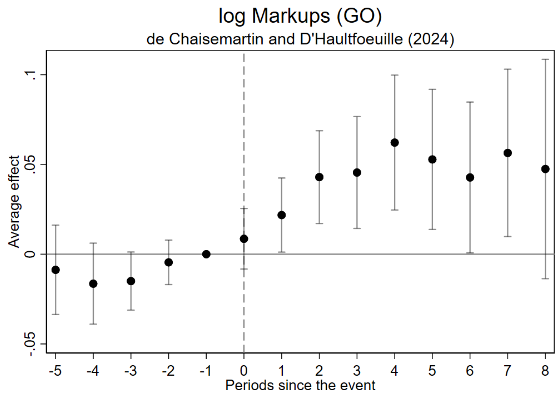

(d) Production-Function Markup

Figure D.1: Effect of M&A on TFPR and Markups (Not-Yet-Treated Control Group)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for the boundary-consistent acquirer–target unit. The control group consists only of not-yet-treated pairs. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used. Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figures 1 and 2, which use the combined control group (never-involved firms and not-yet-treated pairs).
Source: Authors' calculations using Orbis Spain and Orbis M&A.

Table D.1: Event-Study Estimates for Not-Yet-Treated Control Group

	TFPR		Markups	
	Gross Output (1)	Value Added (2)	Inverse Materials Share (3)	Production-Function Markup (4)
Time to Treat −5	0.0169**	0.0279	−0.00646	−0.00871
	(0.00858)	(0.0175)	(0.0137)	(0.0127)
Time to Treat −4	0.0132*	0.0264*	−0.0134	−0.0164
	(0.00761)	(0.0152)	(0.0125)	(0.0115)
Time to Treat −3	−0.00639	0.00183	−0.0217**	−0.0150*
	(0.00642)	(0.0125)	(0.00925)	(0.00826)
Time to Treat −2	−0.00454	−0.00272	−0.00898	−0.00454
	(0.00555)	(0.0116)	(0.00782)	(0.00632)
Time to Treat 0	0.0117*	−0.00223	0.0246**	0.00859
	(0.00609)	(0.0122)	(0.00975)	(0.00861)
Time to Treat 1	−0.00906	0.00231	0.0306***	0.0218**
	(0.00677)	(0.0127)	(0.0117)	(0.0105)
Time to Treat 2	−0.00164	0.0121	0.0570***	0.0429***
	(0.00756)	(0.0151)	(0.0139)	(0.0132)
Time to Treat 3	0.00898	0.0170	0.0652***	0.0455***
	(0.00940)	(0.0171)	(0.0167)	(0.0159)
Time to Treat 4	0.0122	0.0202	0.0782***	0.0622***
	(0.0105)	(0.0180)	(0.0191)	(0.0192)
Time to Treat 5	0.0133	0.0498***	0.0723***	0.0528***
	(0.0103)	(0.0193)	(0.0201)	(0.0199)
Time to Treat 6	0.00672	0.0433*	0.0633***	0.0427**
	(0.0121)	(0.0238)	(0.0222)	(0.0214)
Time to Treat 7	0.0232	0.0883***	0.0854***	0.0564**
	(0.0154)	(0.0284)	(0.0247)	(0.0238)
Time to Treat 8	0.0359**	0.0932***	0.0908***	0.0475
	(0.0177)	(0.0333)	(0.0317)	(0.0312)
ATT	0.00791	0.0238**	0.0549***	0.0370***
	(0.00647)	(0.0119)	(0.0120)	(0.0118)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Event-time estimates from the De Chaisemartin and d'Haultfoeuille (2024) estimator on the pooled acquirer–target unit. The control group consists of not-yet-treated pairs only. Columns (1)–(2) report TFPR from gross-output and value-added production functions. Columns (3)–(4) report the inverse materials share and the production-function markup. Coefficients are log deviations relative to the year before the announcement ($\ell = -1$). Standard errors are clustered at the acquirer–target pair level.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

D.2 Alternative Estimators

The main-text results use the De Chaisemartin and d'Haultfoeuille (2024) estimator, which is designed for staggered settings with heterogeneous treatment effects. We verify that the findings are not driven by the choice of estimator by re-estimating the event studies using four alternatives. Two-way fixed effects (TWFE) is the conventional benchmark, included because it remains the most widely used specification despite its known bias under treatment-effect heterogeneity (De Chaisemartin & d'Haultfoeuille, 2020). Sun and Abraham (2021) correct for contamination from already-treated cohorts by interacting cohort and relative-time indicators. Borusyak et al. (2024) impute untreated potential outcomes for treated units. Callaway and Sant'Anna (2021) estimate group-time average treatment effects and aggregate across cohorts.

Figures D.2–D.5 report the results. The post-merger increases in TFPR and the inverse materials share are present under every estimator. The production-function markup follows the same directional pattern, though estimates are less precise, consistent with the baseline.

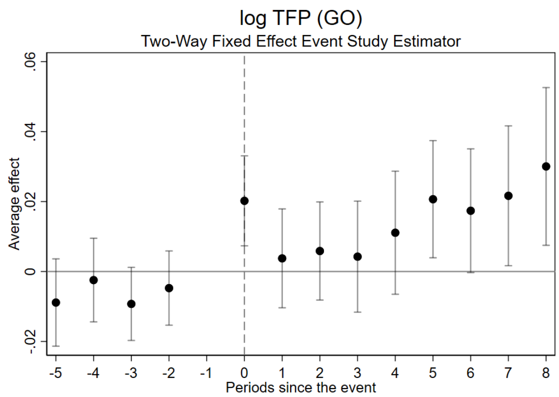

(a) TWFE

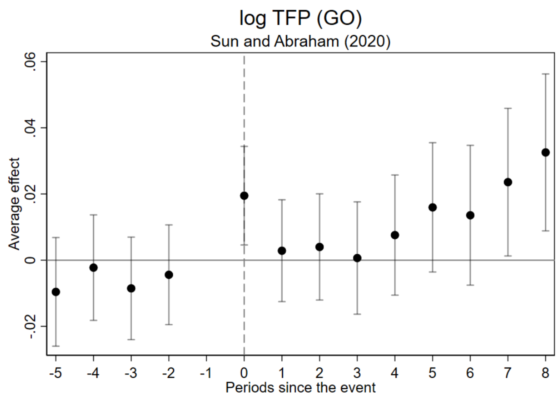

(b) Sun and Abraham (2021)

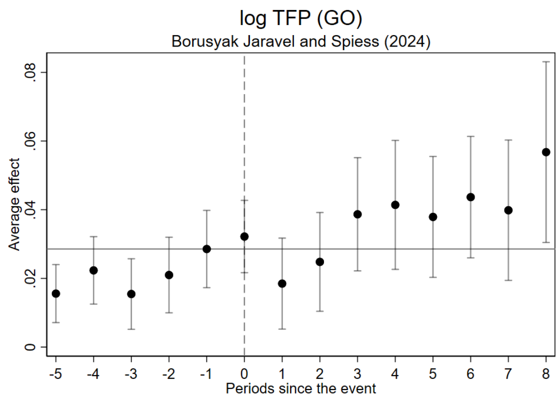

(c) Borusyak et al. (2024)

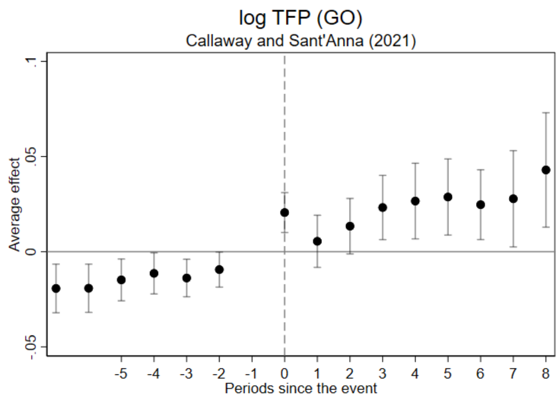

(d) Callaway and Sant'Anna (2021)

Figure D.2: Effect of M&A on Gross-Output TFPR (Alternative Estimators)

Note: Event-study estimates of the effect of horizontal M&A on log gross-output TFPR for the boundary-consistent acquirer–target unit. Each panel reports a different estimator. The control group includes firms never involved in a merger and not-yet-treated pairs for TWFE, Sun and Abraham (2021), and Borusyak et al. (2024). Callaway and Sant'Anna (2021) uses the never-involved control group. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figure 1, Panel A.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

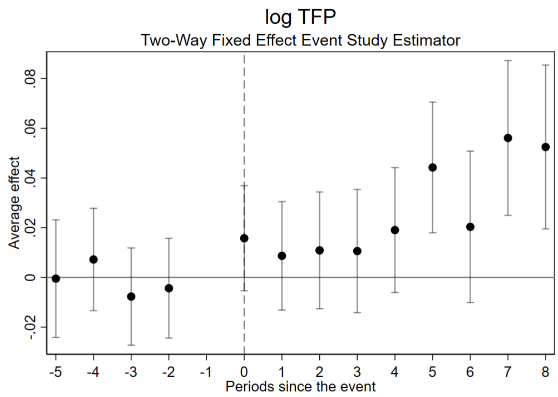

(a) TWFE

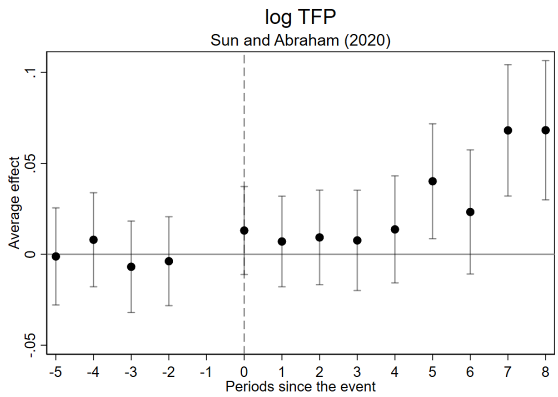

(b) Sun and Abraham (2021)

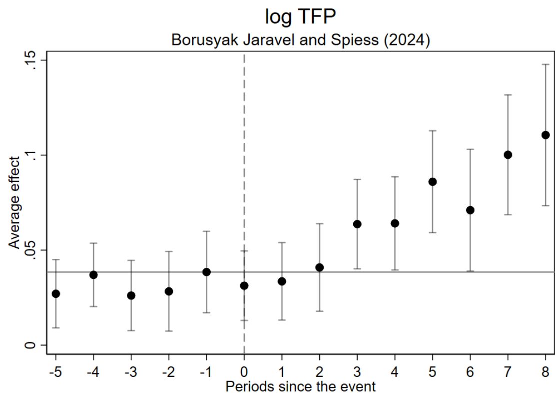

(c) Borusyak et al. (2024)

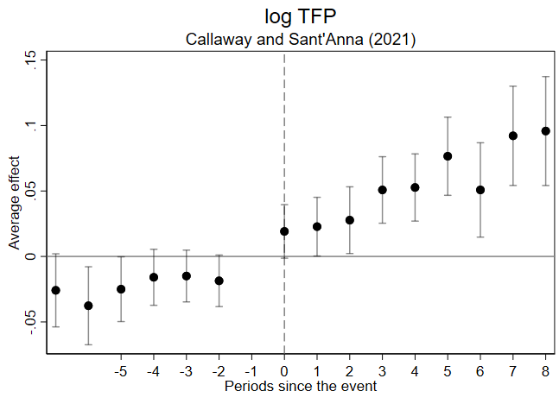

(d) Callaway and Sant'Anna (2021)

Figure D.3: Effect of M&A on Value-Added TFPR (Alternative Estimators)

Note: Event-study estimates of the effect of horizontal M&A on log value-added TFPR for the boundary-consistent acquirer–target unit. Each panel reports a different estimator. The control group includes firms never involved in a merger and not-yet-treated pairs for TWFE, Sun and Abraham (2021), and Borusyak et al. (2024). Callaway and Sant'Anna (2021) uses the never-involved control group. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figure 1, Panel B.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

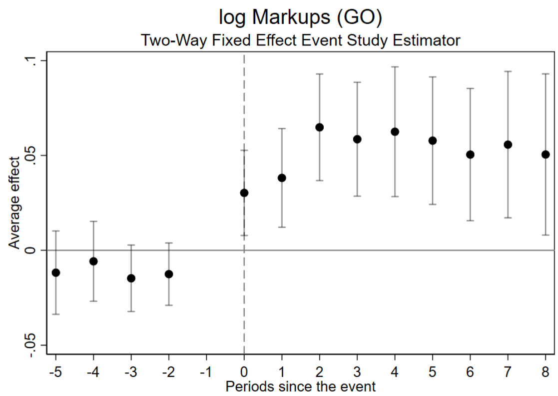

(a) TWFE

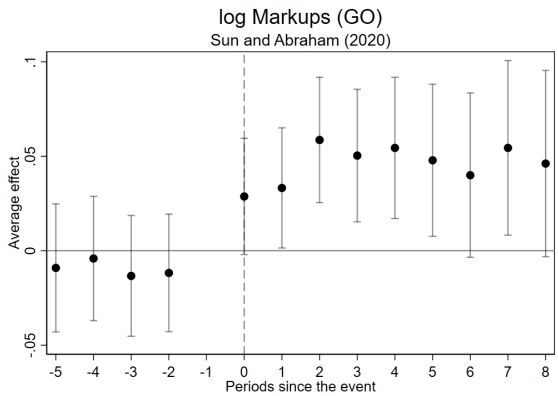

(b) Sun and Abraham (2021)

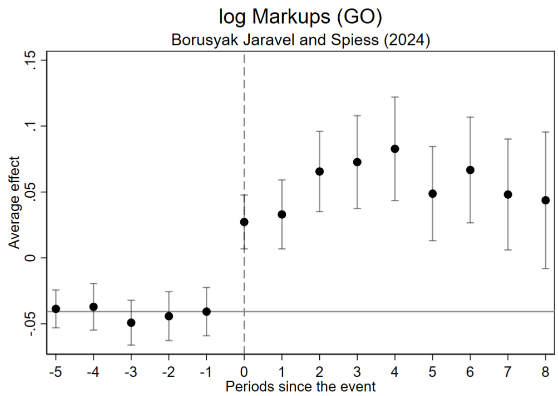

(c) Borusyak et al. (2024)

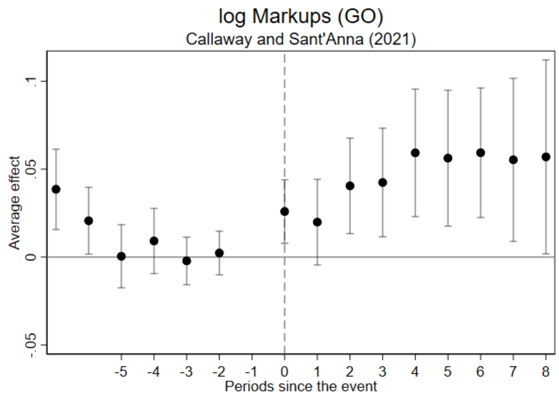

(d) Callaway and Sant'Anna (2021)

Figure D.4: Effect of M&A on Inverse Materials Share (Alternative Estimators)

Note: Event-study estimates of the effect of horizontal M&A on the inverse materials share $\ln(Q_{it}/M_{it})$ for the boundary-consistent acquirer–target unit. Each panel reports a different estimator. The control group includes firms never involved in a merger and not-yet-treated pairs for TWFE, Sun and Abraham (2021), and Borusyak et al. (2024). Callaway and Sant'Anna (2021) uses the never-involved control group. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figure 2, Panel A.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

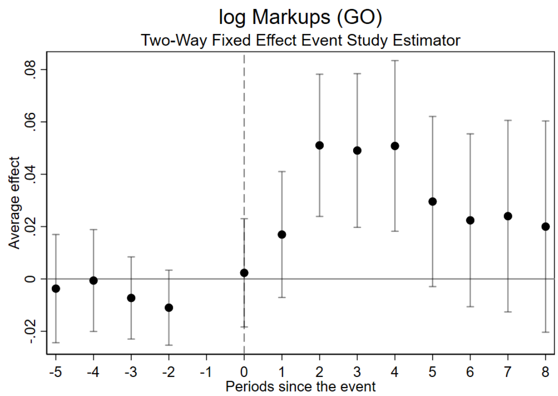

(a) TWFE

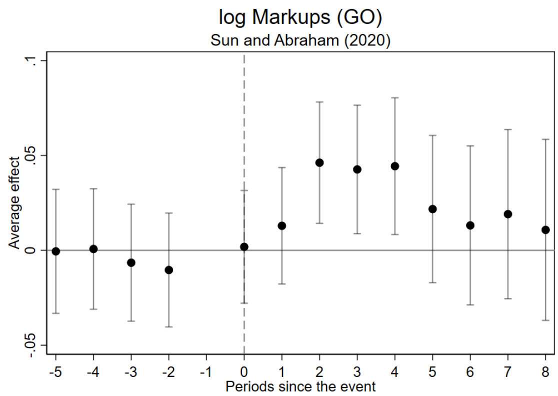

(b) Sun and Abraham (2021)

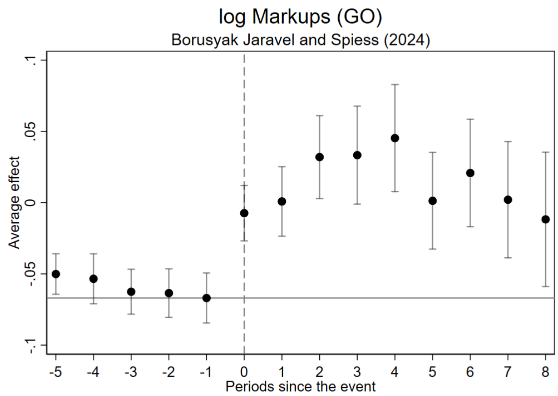

(c) Borusyak et al. (2024)

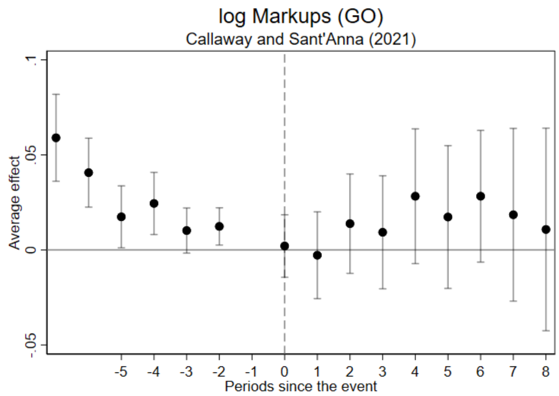

(d) Callaway and Sant'Anna (2021)

Figure D.5: Effect of M&A on Production-Function Markup (Alternative Estimators)

Note: Event-study estimates of the effect of horizontal M&A on the production-function markup for the boundary-consistent acquirer–target unit. The markup is estimated using fitted gross output to address measurement error. Each panel reports a different estimator. The control group includes firms never involved in a merger and not-yet-treated pairs for TWFE, Sun and Abraham (2021), and Borusyak et al. (2024). Callaway and Sant'Anna (2021) uses the never-involved control group. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figure 2, Panel B.

Source: Authors' calculations using Orbis Spain and Orbis M&A.

Appendix E Industry-Level Event Studies

Table 2 in the main text summarizes the industry-level TFPR and markup patterns for the seven industries with reliable pre-trends. This appendix reports the full event-study dynamics for all ten NACE 2-digit industries with the largest shares of horizontal merger activity, ordered by their share of total deals. Seven of the ten industries satisfy the parallel-trends assumption and show stable post-merger dynamics across all four outcomes: wholesale trade (NACE 46), retail trade (NACE 47), motor vehicle trade (NACE 45), land transport (NACE 49), human health activities (NACE 86), food manufacturing (NACE 10), and specialized construction (NACE 43). The remaining three — construction of buildings (NACE 41), computer programming (NACE 62), and fabricated metal products (NACE 25) — show pre-trend violations or implausible post-event dynamics that make their estimates difficult to interpret.

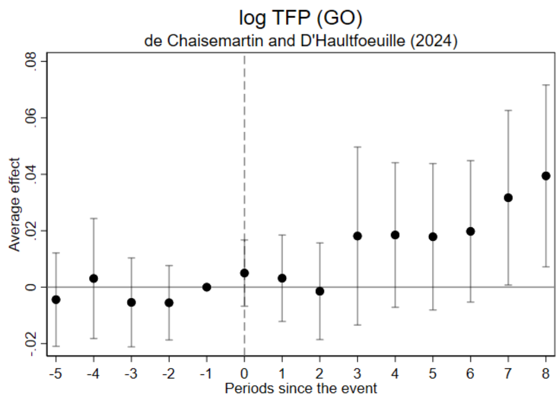

(a) Gross-Output TFPR

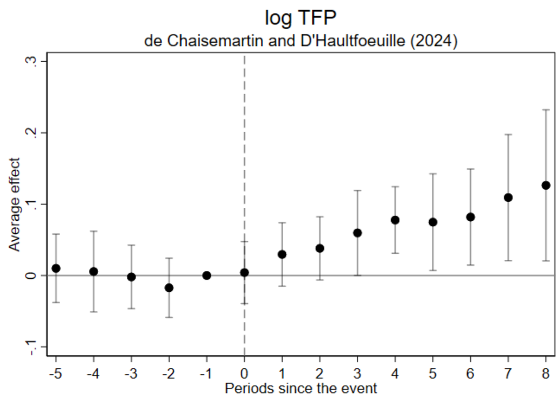

(b) Value-Added TFPR

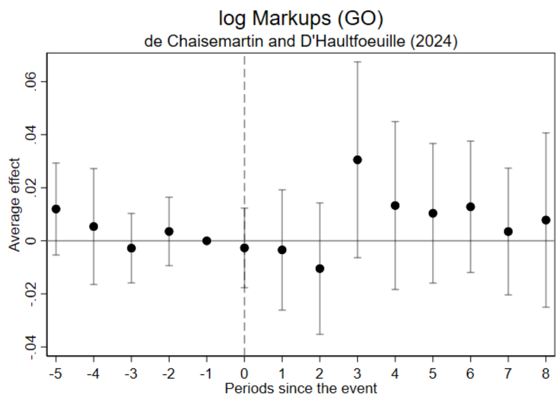

(c) Inverse Materials Share

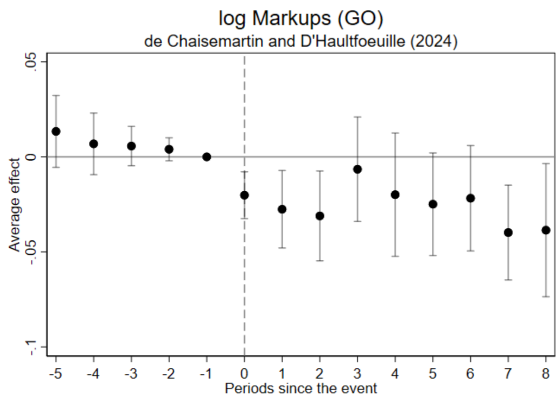

(d) Production-Function Markup

Figure E.1: NACE 46 (Wholesale Trade)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 46 (Wholesale trade, except of motor vehicles and motorcycles). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 18.3% of horizontal deals.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

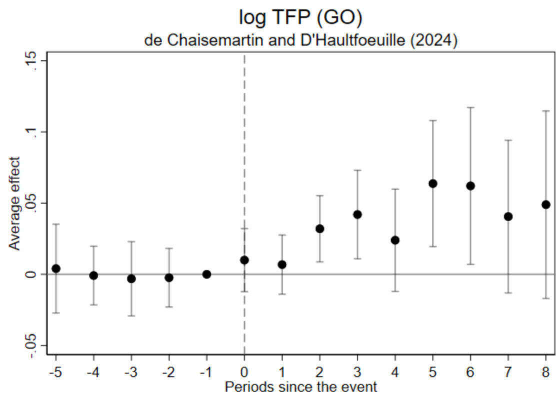

(a) Gross-Output TFPR

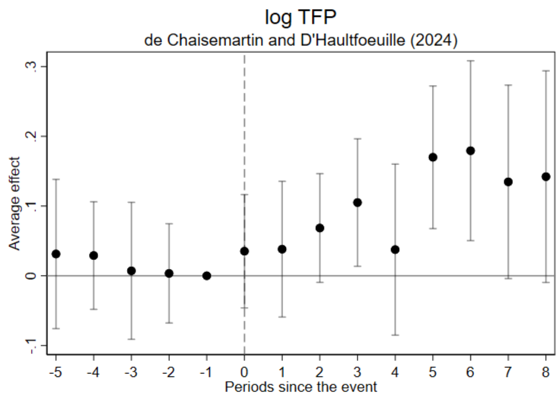

(b) Value-Added TFPR

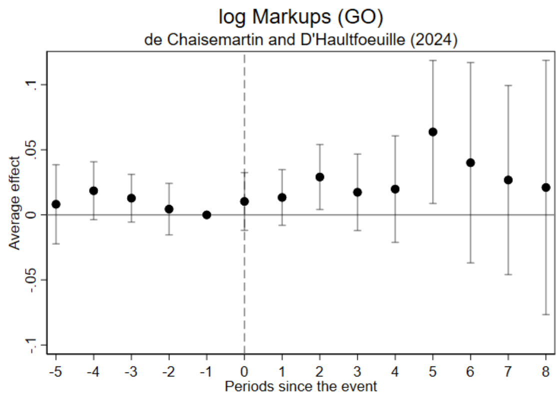

(c) Inverse Materials Share

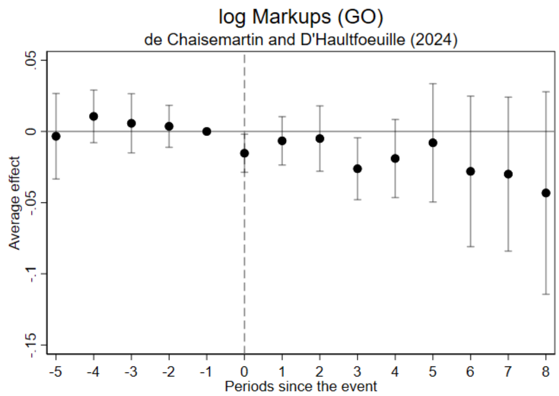

(d) Production-Function Markup

Figure E.2: NACE 47 (Retail Trade)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 47 (Retail trade, except of motor vehicles and motorcycles). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 5.0% of horizontal deals.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

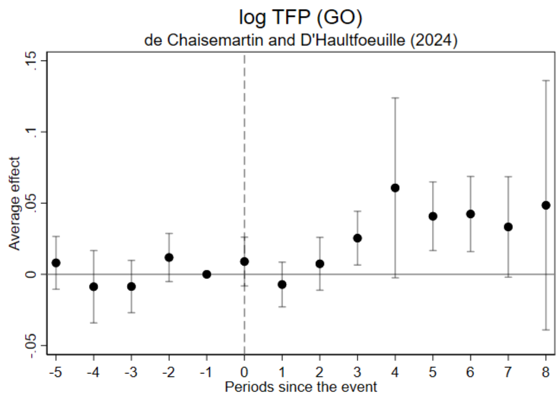

(a) Gross-Output TFPR

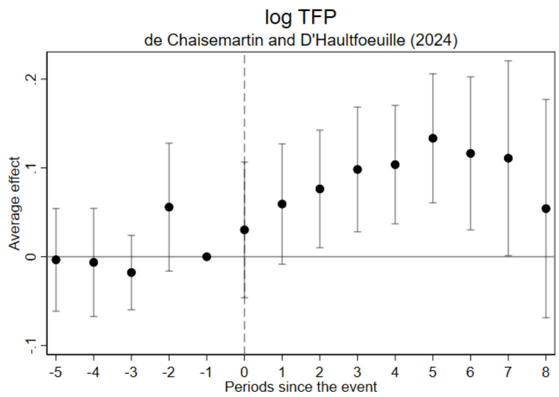

(b) Value-Added TFPR

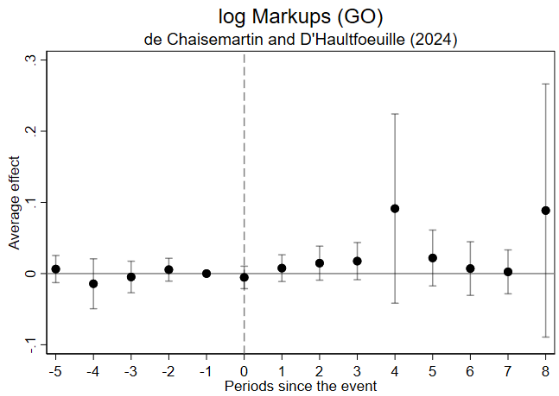

(c) Inverse Materials Share

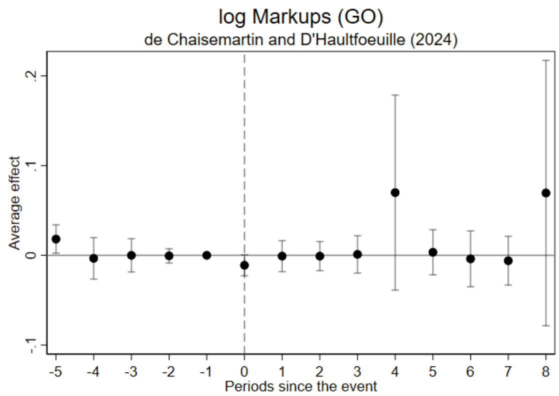

(d) Production-Function Markup

Figure E.3: NACE 45 (Motor Vehicle Trade)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 45 (Wholesale and retail trade and repair of motor vehicles and motorcycles). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 4.9% of horizontal deals.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

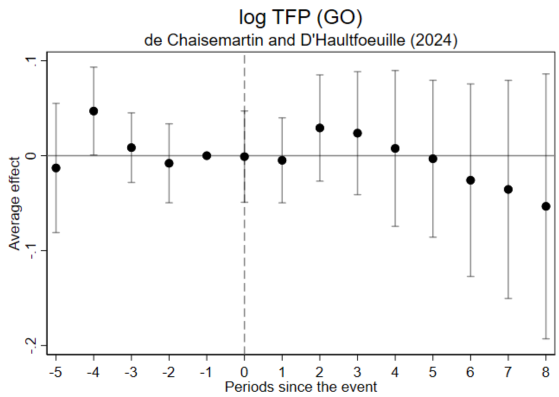

(a) Gross-Output TFPR

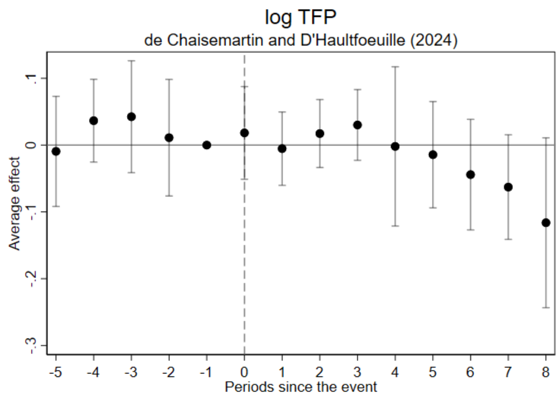

(b) Value-Added TFPR

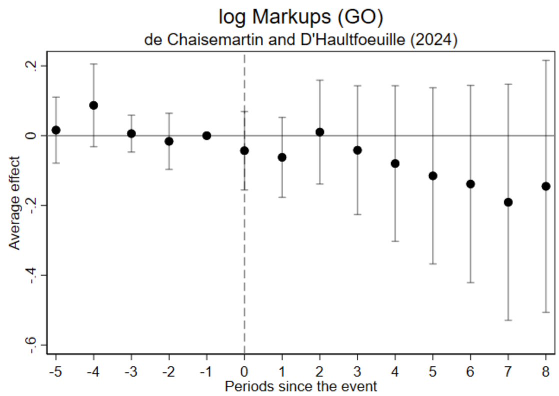

(c) Inverse Materials Share

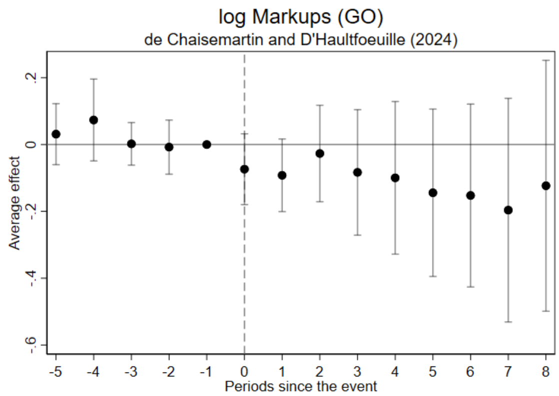

(d) Production-Function Markup

Figure E.4: NACE 49 (Land Transport)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 49 (Land transport and transport via pipelines). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 4.4% of horizontal deals.

Source: Authors' calculations using Orbis Spain and Orbis M&A.

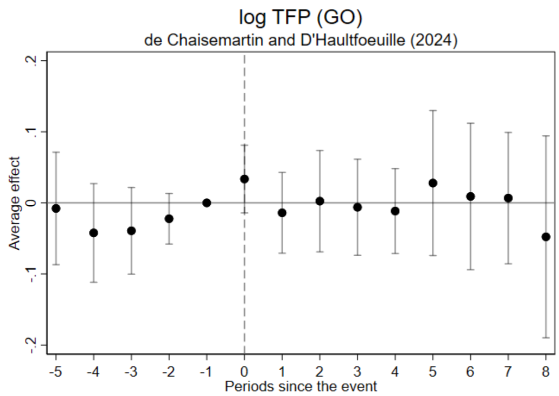

(a) Gross-Output TFPR

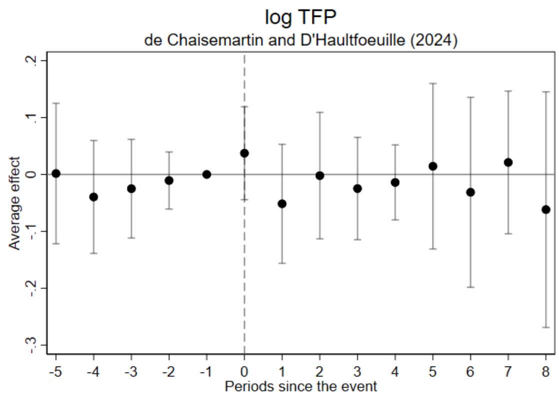

(b) Value-Added TFPR

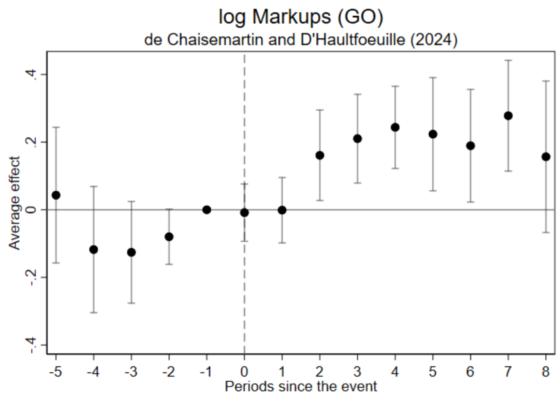

(c) Inverse Materials Share

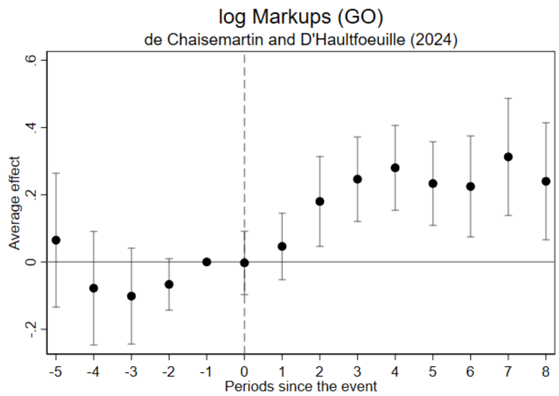

(d) Production-Function Markup

Figure E.5: NACE 86 (Human Health Activities)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 86 (Human health activities). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 3.5% of horizontal deals.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

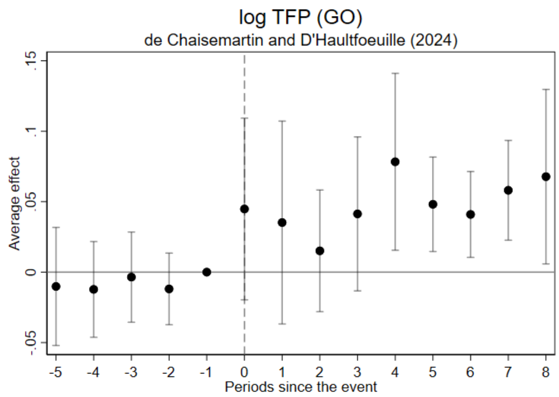

(a) Gross-Output TFPR

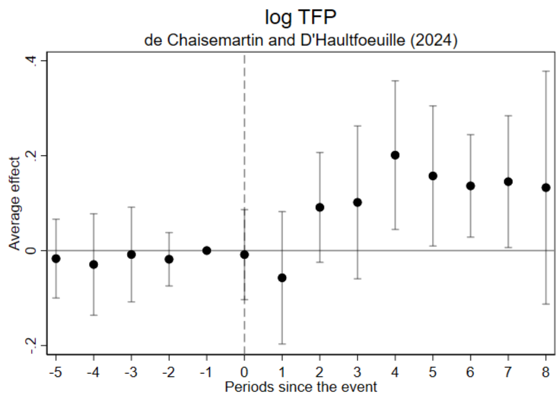

(b) Value-Added TFPR

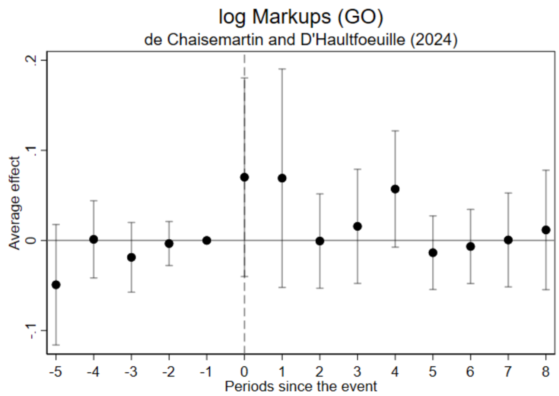

(c) Inverse Materials Share

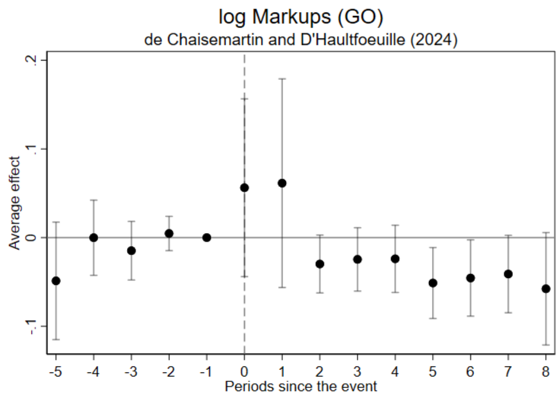

(d) Production-Function Markup

Figure E.6: NACE 10 (Food Manufacturing)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 10 (Manufacture of food products). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 3.4% of horizontal deals.

Source: Authors' calculations using Orbis Spain and Orbis M&A.

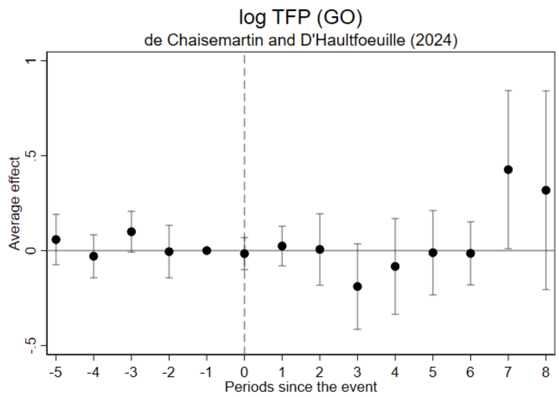

(a) Gross-Output TFPR

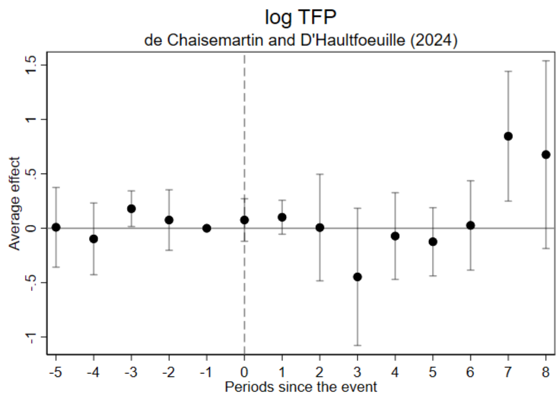

(b) Value-Added TFPR

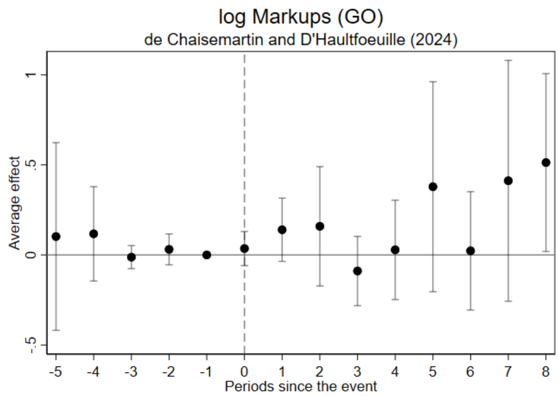

(c) Inverse Materials Share

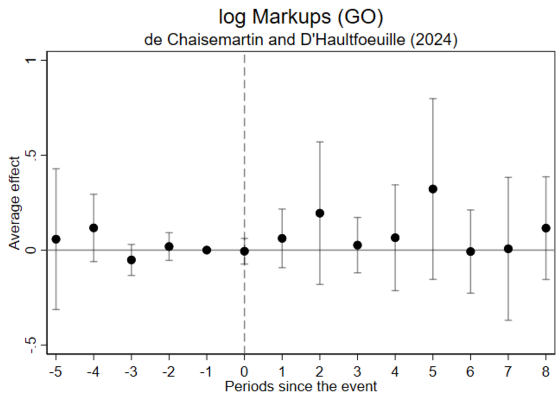

(d) Production-Function Markup

Figure E.7: NACE 41 (Construction of Buildings)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 41 (Construction of buildings). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 3.3% of horizontal deals. *This industry exhibits late-horizon spikes inconsistent with gradual treatment effects; post-merger estimates should be interpreted with caution.*
Source: Authors' calculations using Orbis Spain and Orbis M&A.

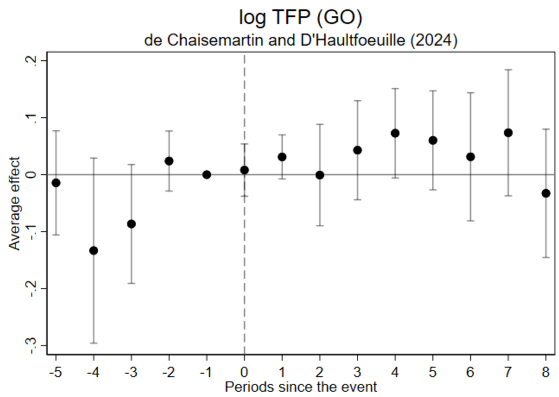

(a) Gross-Output TFPR

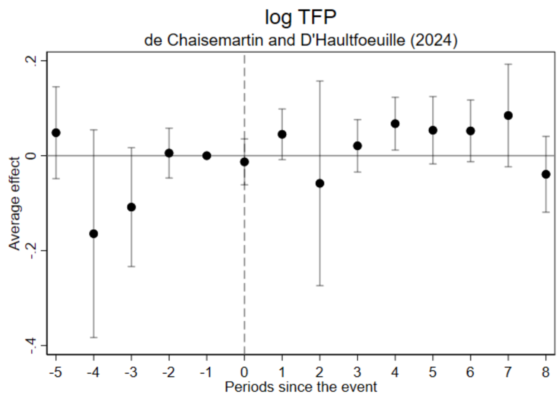

(b) Value-Added TFPR

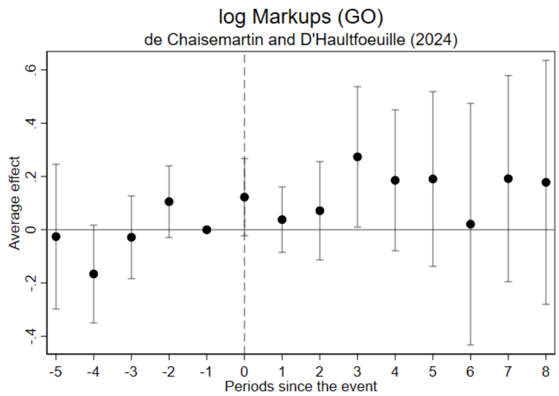

(c) Inverse Materials Share

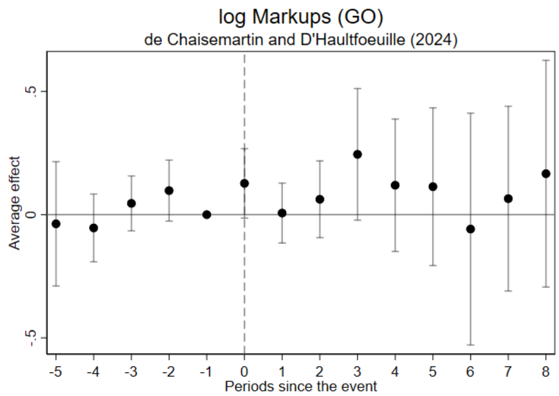

(d) Production-Function Markup

Figure E.8: NACE 62 (Computer Programming)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 62 (Computer programming, consultancy, and related activities). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 3.2% of horizontal deals. *This industry shows TFPR pre-trend violations; post-merger estimates should be interpreted with caution.*

Source: Authors' calculations using Orbis Spain and Orbis M&A.

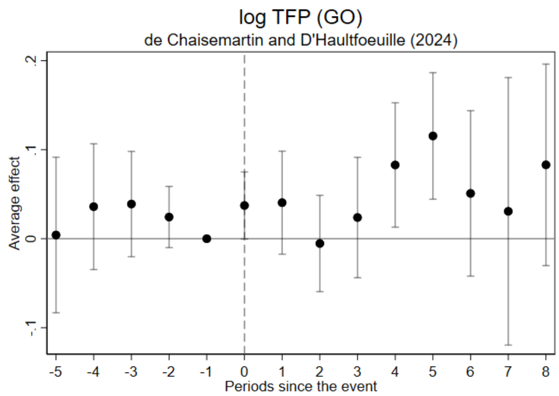

(a) Gross-Output TFPR

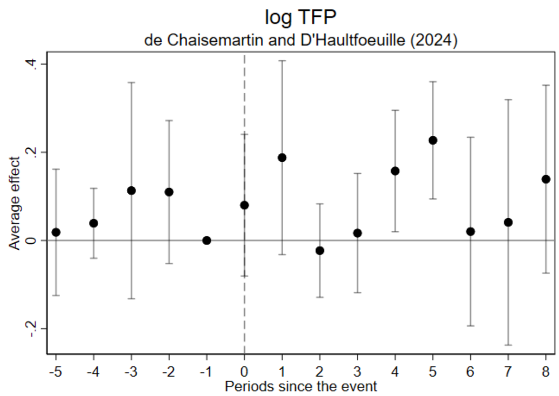

(b) Value-Added TFPR

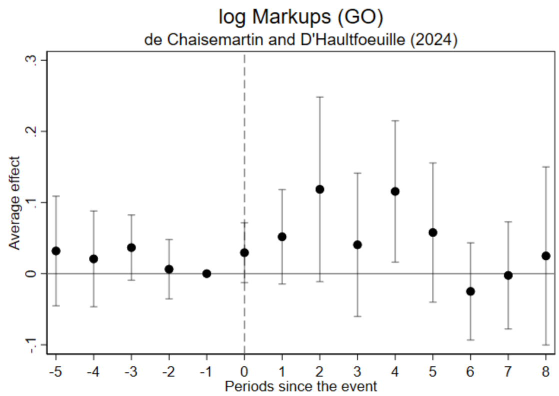

(c) Inverse Materials Share

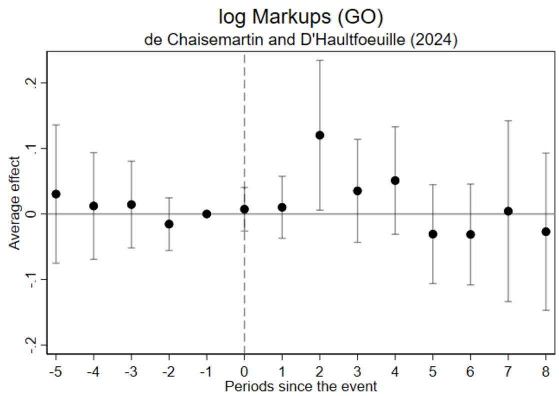

(d) Production-Function Markup

Figure E.9: NACE 25 (Fabricated Metal Products)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 25 (Manufacture of fabricated metal products, except machinery and equipment). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 2.6% of horizontal deals. *This industry shows value-added TFPR pre-trend violations; post-merger estimates should be interpreted with caution.*

Source: Authors' calculations using Orbis Spain and Orbis M&A.

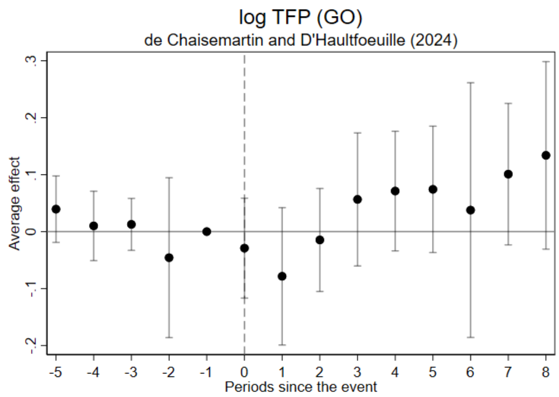

(a) Gross-Output TFPR

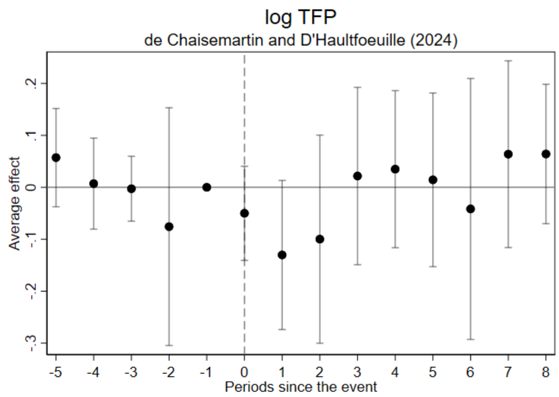

(b) Value-Added TFPR

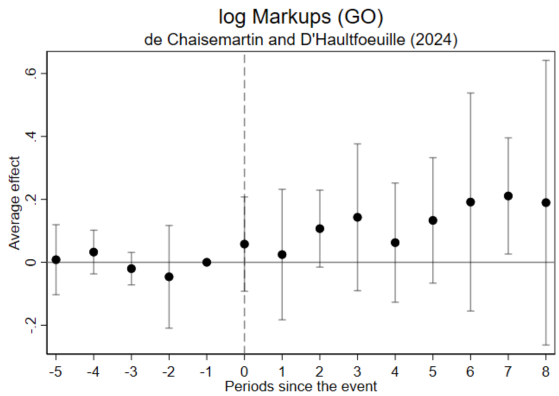

(c) Inverse Materials Share

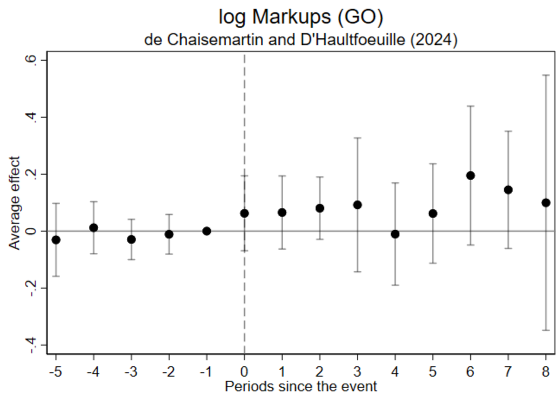

(d) Production-Function Markup

Figure E.10: NACE 43 (Specialized Construction)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for NACE 43 (Specialized construction activities). Outcomes are measured on the boundary-consistent acquirer–target unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. This industry accounts for 2.6% of horizontal deals.

Source: Authors' calculations using Orbis Spain and Orbis M&A.

Appendix F Robustness

This appendix section reports figures supporting the robustness checks discussed in Section 5.

F.1 Matched Difference-in-Differences

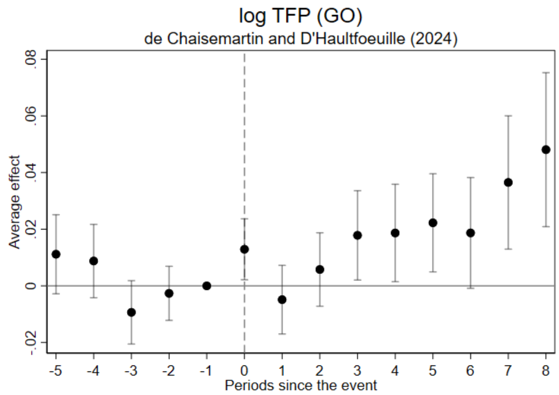

(a) Gross-Output TFPR

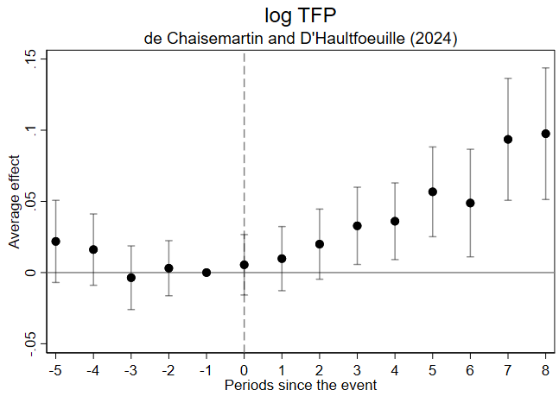

(b) Value-Added TFPR

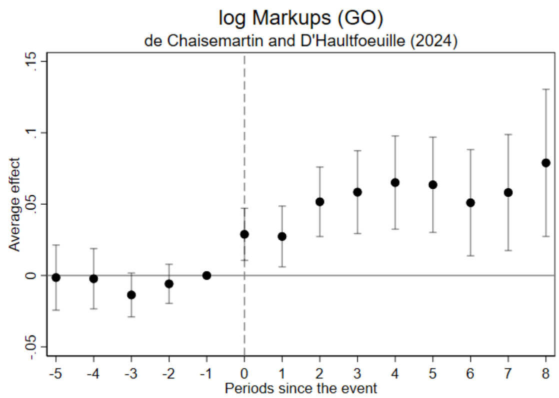

(c) Inverse Materials Share

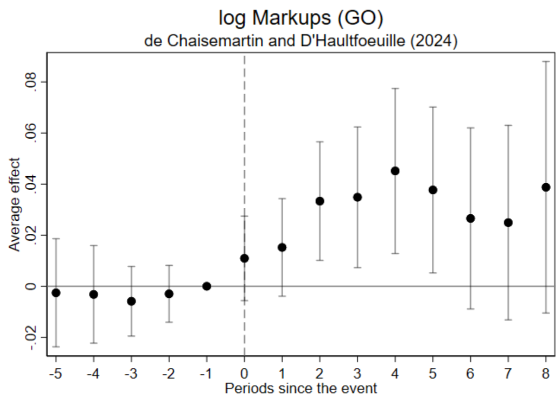

(d) Production-Function Markup

Figure F.1: Effect of M&A on TFPR and Markups (Propensity-Score-Matched Sample)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for the boundary-consistent acquirer–target unit, estimated on a propensity-score-matched sample. Each treated pair is matched to one control firm using nearest-neighbor matching (caliper of 0.01) on pre-treatment log total assets, tangible fixed assets, labor costs, material costs, and sales. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figures 1 and 2, which report the baseline specification on the unmatched sample.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

F.2 Four-Digit Industry Classification

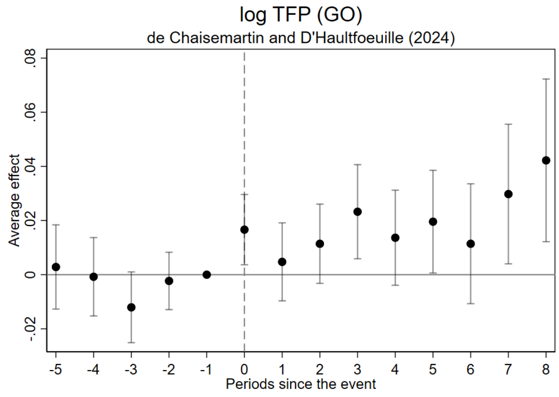

(a) Gross-Output TFPR

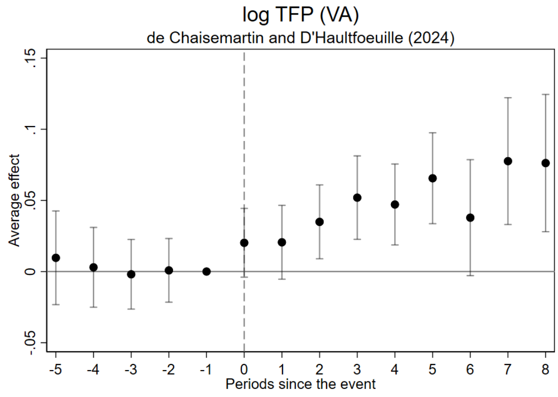

(b) Value-Added TFPR

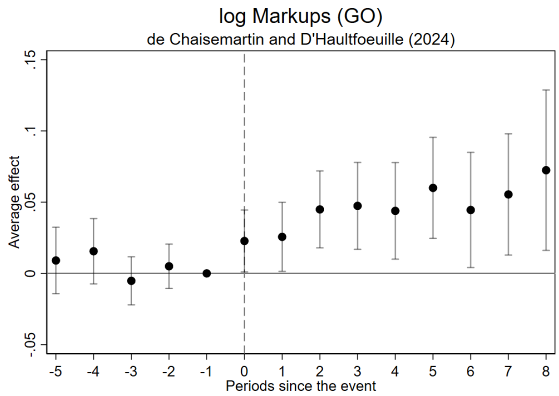

(c) Inverse Materials Share

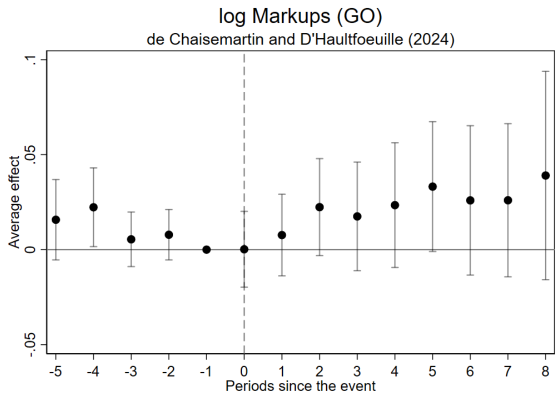

(d) Production-Function Markup

Figure F.2: Effect of M&A on TFPR and Markups (NACE 4-Digit Horizontal Definition)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for the boundary-consistent acquirer–target unit. Horizontal mergers are defined as those in which the acquirer and target share the same NACE 4-digit industry, reducing the sample from 3,759 to 2,809 deals. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figures 1 and 2, which use the NACE 2-digit definition.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

F.3 Alternative Output and Capital Measures

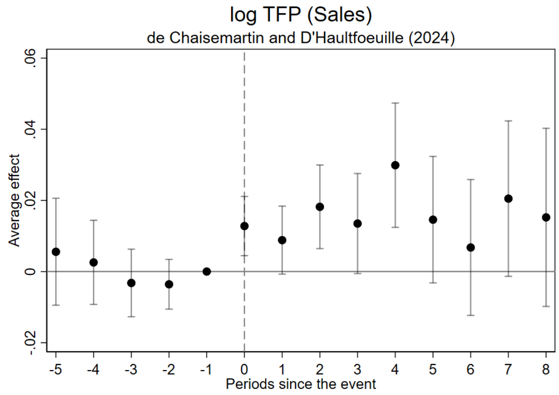

(a) Sales TFPR

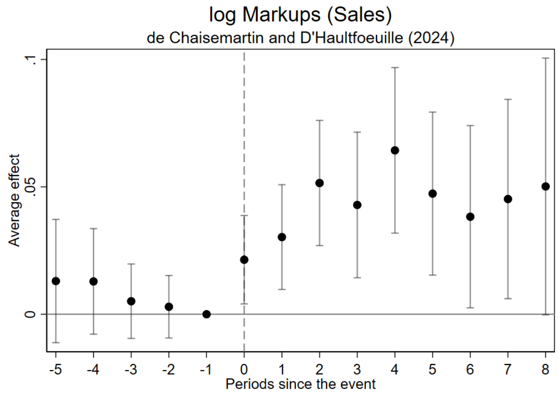

(b) Inverse Materials Share (Sales)

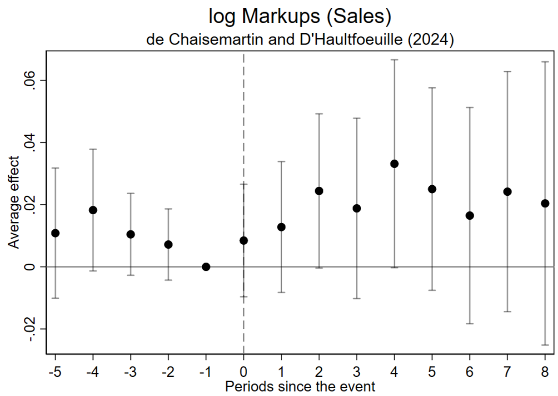

(c) Production-Function Markup (Sales)

Figure F.3: Effect of M&A on TFPR and Markups (Sales as Output Measure)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panel A) and markups (Panels B–C) for the boundary-consistent acquirer–target unit, using sales as the output measure in place of gross output. The inverse materials share (Panel B) is computed as $\ln(\text{Sales}_{it}/M_{it})$. The production-function markup (Panel C) uses fitted sales from the estimated production function. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figures 1 and 2, which use gross output.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

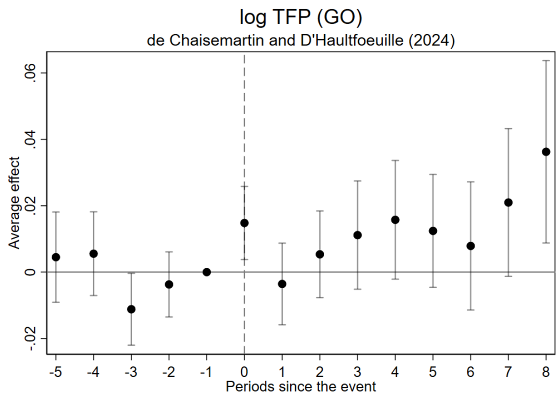

(a) Gross-Output TFPR

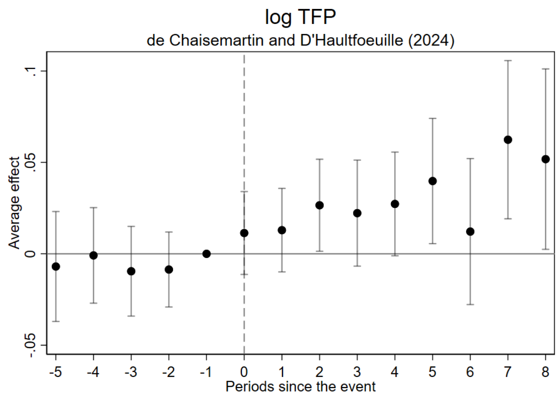

(b) Value-Added TFPR

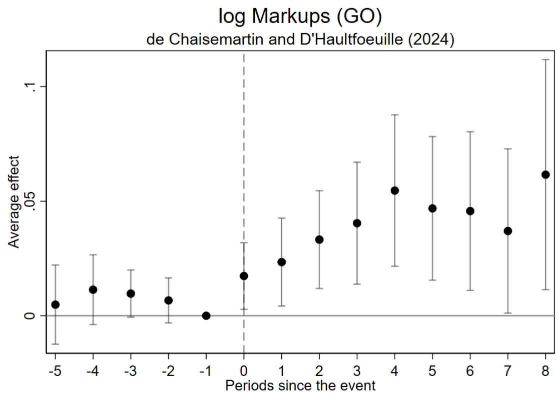

(c) Production-Function Markup

Figure F.4: Effect of M&A on TFPR and Markups (Tangible Fixed Assets as Capital Proxy)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panel C) for the boundary-consistent acquirer–target unit, using tangible fixed assets as the capital proxy in place of total assets. The inverse materials share is not affected by the capital proxy and is therefore omitted. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figures 1 and 2, which use total assets.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

F.4 Inclusion of Cross-Industry Mergers

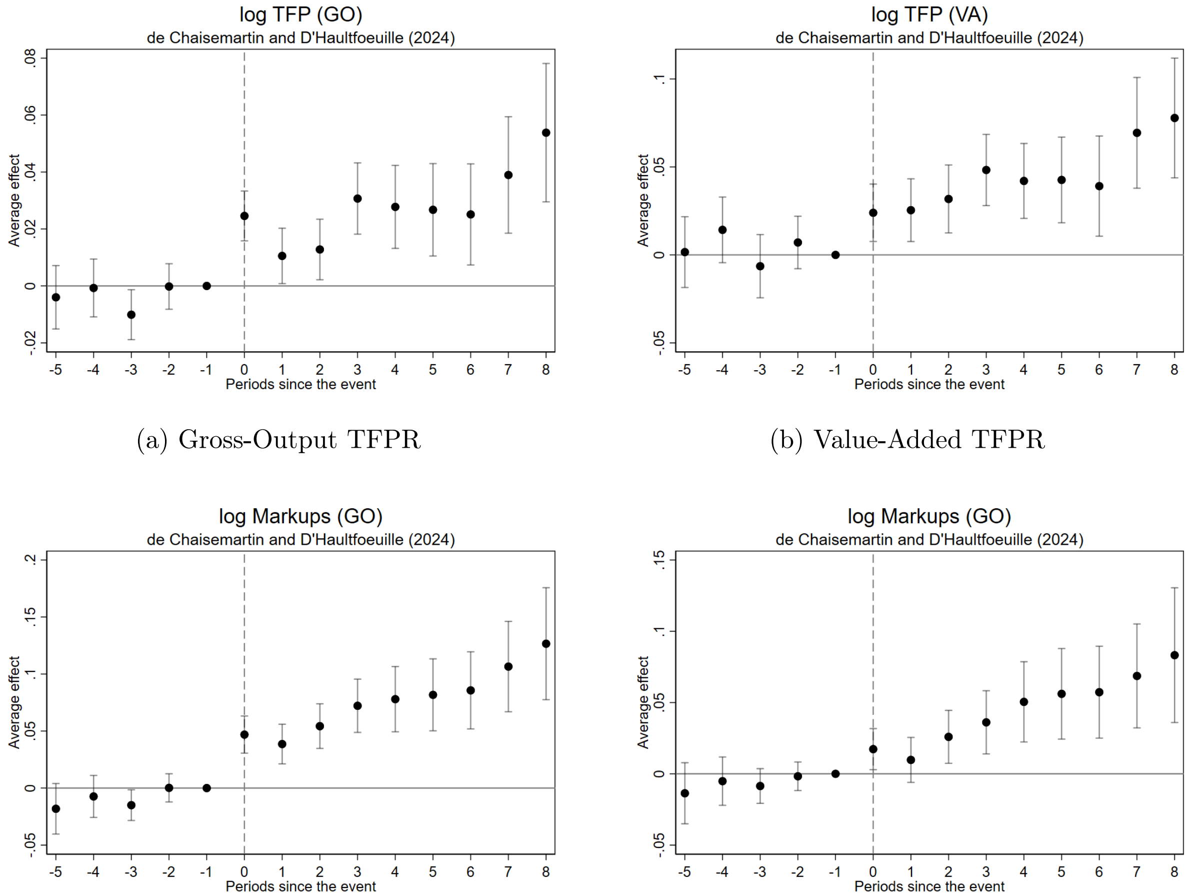

(a) Gross-Output TFPR

(b) Value-Added TFPR

(c) Inverse Materials Share

(d) Production-Function Markup

Figure F.5: Effect of M&A on TFPR and Markups (Including Cross-Industry Mergers)

Note: Event-study estimates of the effect of M&A on TFPR (Panels A–B) and markups (Panels C–D) for the boundary-consistent acquirer–target unit, including both horizontal and cross-industry mergers. For cross-industry deals, the acquirer's industry-level output elasticities are assigned to the combined unit. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figures 1 and 2, which restrict to horizontal mergers (same NACE 2-digit industry).
Source: Authors' calculations using Orbis Spain and Orbis M&A.

F.5 Excluding Misrecorded Deal Dates

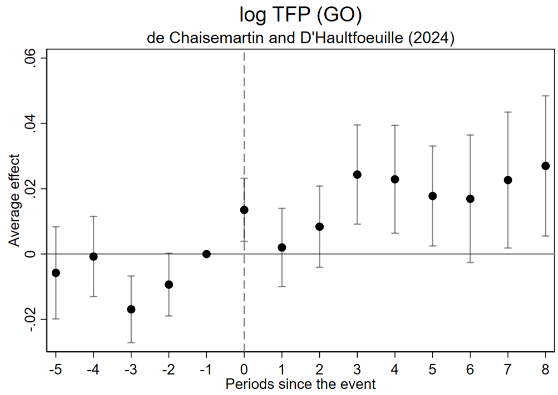

(a) Gross-Output TFPR

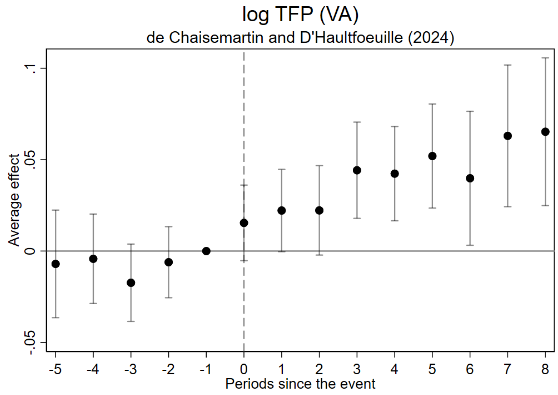

(b) Value-Added TFPR

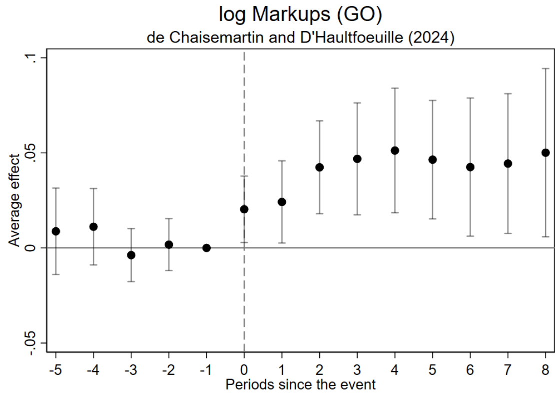

(c) Inverse Materials Share

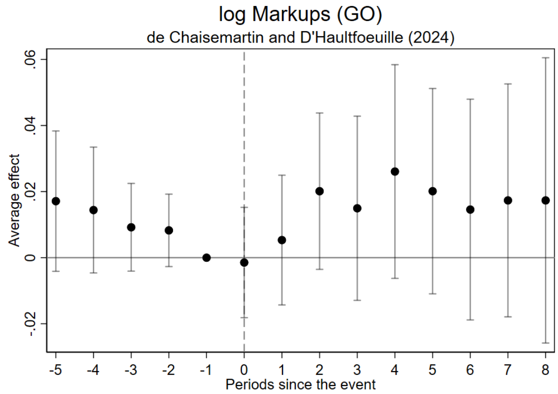

(d) Production-Function Markup

Figure F.6: Effect of M&A on TFPR and Markups (Dropping Deals with Two-Year Date Gap)

Note: Event-study estimates of the effect of horizontal M&A on TFPR (Panels A–B) and markups (Panels C–D) for the boundary-consistent acquirer–target unit. Deals in which the target disappears two or more years before the recorded announcement date are dropped rather than corrected. The De Chaisemartin and d'Haultfoeuille (2024) estimator is used with the combined control group (never-involved firms and not-yet-treated pairs). Each point is an event-time ATT relative to $\ell = -1$. Bars are 95% confidence intervals with standard errors clustered at the pair level. Compare to Figures 1 and 2, which reassign the deal year for these cases.
Source: Authors' calculations using Orbis Spain and Orbis M&A.

Appendix G Formal Results and Proofs

G.1 Proposition 1 (Consolidation bias)

We restate the proposition here with all notation defined.

Proposition 1 *(Consolidation bias identity and threshold characterization).* Suppose two production units $i \in \{A, T\}$ satisfy the Cobb–Douglas gross-output revenue technology

$$Q_i = Z_i K_i^{\beta_k} L_i^{\beta_l} M_i^{\beta_m}, \qquad \beta_k, \beta_l, \beta_m > 0, \qquad K_i > 0,\ L_i > 0,\ M_i > 0,\ Z_i > 0.$$

Define consolidated ("pooled") accounts

$$Q^{\text{pooled}} \equiv Q_A + Q_T, \qquad K^{\text{pooled}} \equiv K_A + K_T, \qquad L^{\text{pooled}} \equiv L_A + L_T, \qquad M^{\text{pooled}} \equiv M_A + M_T,$$

and the TFPR computed from consolidated accounts,

$$Z^{\text{pooled}} \equiv \frac{Q^{\text{pooled}}}{(K^{\text{pooled}})^{\beta_k}(L^{\text{pooled}})^{\beta_l}(M^{\text{pooled}})^{\beta_m}}.$$

Let

$$s_K \equiv \frac{K_A}{K_A + K_T} \in (0,1), \qquad s_L \equiv \frac{L_A}{L_A + L_T} \in (0,1), \qquad s_M \equiv \frac{M_A}{M_A + M_T} \in (0,1), \qquad r \equiv \frac{Z_T}{Z_A} > 0,$$

and define

$$a \equiv s_K^{\beta_k} s_L^{\beta_l} s_M^{\beta_m}, \qquad b \equiv (1-s_K)^{\beta_k}(1-s_L)^{\beta_l}(1-s_M)^{\beta_m}.$$

Then:

(i) (**Identity**) The consolidated-to-acquirer TFPR ratio satisfies

$$\frac{Z^{\text{pooled}}}{Z_A} = a + rb. \tag{G.1}$$

Equivalently, $\ln Z^{\text{pooled}} - \ln Z_A = \ln(a + rb)$.

(ii) (**Knife-edge equality**) Define

$$r^* \equiv \frac{1-a}{b} = \frac{1 - s_K^{\beta_k} s_L^{\beta_l} s_M^{\beta_m}}{(1-s_K)^{\beta_k}(1-s_L)^{\beta_l}(1-s_M)^{\beta_m}}. \tag{G.2}$$

Then $r^* > 0$ and $Z^{\text{pooled}} = Z_A \iff r = r^*$.

(iii) (**Sign characterization**) With r^* as above,

$$\ln Z^{\text{pooled}} - \ln Z_A \gtrless 0 \quad \iff \quad r \gtrless r^*.$$

Without restrictions on the target's relative productivity r and the within-consolidated input shares (s_K, s_L, s_M), the sign of consolidation bias is not identified from consolidated accounts alone.

Proof. We proceed in three steps corresponding to parts (i)–(iii).

Step 1: Proof of (i) (Identity). By definition,

$$Z^{\text{pooled}} = \frac{Q^{\text{pooled}}}{(K^{\text{pooled}})^{\beta_k}(L^{\text{pooled}})^{\beta_l}(M^{\text{pooled}})^{\beta_m}} = \frac{Q_A + Q_T}{(K_A + K_T)^{\beta_k}(L_A + L_T)^{\beta_l}(M_A + M_T)^{\beta_m}}.$$

Substituting $Q_i = Z_i K_i^{\beta_k} L_i^{\beta_l} M_i^{\beta_m}$ for $i \in \{A, T\}$ yields

$$Z^{\text{pooled}} = \frac{Z_A K_A^{\beta_k} L_A^{\beta_l} M_A^{\beta_m} + Z_T K_T^{\beta_k} L_T^{\beta_l} M_T^{\beta_m}}{(K_A + K_T)^{\beta_k}(L_A + L_T)^{\beta_l}(M_A + M_T)^{\beta_m}}.$$

Since $Z_A > 0$, divide both sides by Z_A:

$$\frac{Z^{\text{pooled}}}{Z_A} = \frac{K_A^{\beta_k} L_A^{\beta_l} M_A^{\beta_m}}{(K_A + K_T)^{\beta_k}(L_A + L_T)^{\beta_l}(M_A + M_T)^{\beta_m}} + \frac{Z_T}{Z_A} \cdot \frac{K_T^{\beta_k} L_T^{\beta_l} M_T^{\beta_m}}{(K_A + K_T)^{\beta_k}(L_A + L_T)^{\beta_l}(M_A + M_T)^{\beta_m}}. \tag{G.3}$$

Now define

$$s_K \equiv \frac{K_A}{K_A + K_T}, \qquad s_L \equiv \frac{L_A}{L_A + L_T}, \qquad s_M \equiv \frac{M_A}{M_A + M_T}.$$

Because $K_A, K_T > 0$, $L_A, L_T > 0$, and $M_A, M_T > 0$, we have $K_A + K_T > 0$, $L_A + L_T > 0$, and $M_A + M_T > 0$, hence $s_K, s_L, s_M \in (0, 1)$. Moreover,

$$K_A = s_K(K_A + K_T), \qquad K_T = (1 - s_K)(K_A + K_T),$$

$$L_A = s_L(L_A + L_T), \qquad L_T = (1 - s_L)(L_A + L_T),$$

$$M_A = s_M(M_A + M_T), \qquad M_T = (1 - s_M)(M_A + M_T).$$

Substitute these into the first fraction on the right-hand side of (G.3):

$$\begin{aligned} \frac{K_A^{\beta_k} L_A^{\beta_l} M_A^{\beta_m}}{(K_A + K_T)^{\beta_k}(L_A + L_T)^{\beta_l}(M_A + M_T)^{\beta_m}} &= \frac{\big(s_K(K_A + K_T)\big)^{\beta_k}\big(s_L(L_A + L_T)\big)^{\beta_l}\big(s_M(M_A + M_T)\big)^{\beta_m}}{(K_A + K_T)^{\beta_k}(L_A + L_T)^{\beta_l}(M_A + M_T)^{\beta_m}} \\ &= s_K^{\beta_k} s_L^{\beta_l} s_M^{\beta_m}. \end{aligned}$$

Likewise, the second fraction becomes

$$\begin{aligned} \frac{K_T^{\beta_k} L_T^{\beta_l} M_T^{\beta_m}}{(K_A + K_T)^{\beta_k}(L_A + L_T)^{\beta_l}(M_A + M_T)^{\beta_m}} &= \frac{\big((1 - s_K)(K_A + K_T)\big)^{\beta_k}\big((1 - s_L)(L_A + L_T)\big)^{\beta_l}\big((1 - s_M)(M_A + M}{(K_A + K_T)^{\beta_k}(L_A + L_T)^{\beta_l}(M_A + M_T)^{\beta_m}} \\ &= (1 - s_K)^{\beta_k}(1 - s_L)^{\beta_l}(1 - s_M)^{\beta_m}. \end{aligned}$$

Define $r \equiv Z_T / Z_A > 0$,

$$a \equiv s_K^{\beta_k} s_L^{\beta_l} s_M^{\beta_m}, \qquad b \equiv (1 - s_K)^{\beta_k}(1 - s_L)^{\beta_l}(1 - s_M)^{\beta_m}.$$

Substituting into (G.3) yields

$$\frac{Z^{\text{pooled}}}{Z_A} = a + rb,$$

which is (G.1). For the log representation, note that $Z^{\text{pooled}} > 0$ and $Z_A > 0$, so the ratio Z^{pooled}/Z_A is positive and $\ln(\cdot)$ is well-defined. Using $\ln x - \ln y = \ln(x/y)$ for $x, y > 0$ gives

$$\ln Z^{\text{pooled}} - \ln Z_A = \ln\left(\frac{Z^{\text{pooled}}}{Z_A}\right) = \ln(a + rb).$$

Step 2: Proof of (ii) (Knife-edge equality). First show $a \in (0,1)$ and $b \in (0,1)$. Since $s_K, s_L, s_M \in (0,1)$ and $\beta_k, \beta_l, \beta_m > 0$, we have $s_K^{\beta_k}, s_L^{\beta_l}, s_M^{\beta_m} \in (0,1)$, hence $a = s_K^{\beta_k} s_L^{\beta_l} s_M^{\beta_m} \in (0,1)$. Similarly, $(1-s_K), (1-s_L), (1-s_M) \in (0,1)$ implies $b = (1-s_K)^{\beta_k}(1-s_L)^{\beta_l}(1-s_M)^{\beta_m} \in (0,1)$, in particular $b > 0$.

Define $r^* = (1-a)/b$ as in (G.2). Since $a \in (0,1)$ implies $1-a > 0$ and $b > 0$, it follows that $r^* > 0$.

We now prove the equivalence $Z^{\text{pooled}} = Z_A \iff r = r^*$.

($\Rightarrow$) Assume $Z^{\text{pooled}} = Z_A$. Because $Z_A > 0$, dividing both sides by Z_A gives $Z^{\text{pooled}}/Z_A = 1$. By part (i), $Z^{\text{pooled}}/Z_A = a + rb$, hence

$$1 = a + rb.$$

Since $b > 0$, rearranging yields

$$r = \frac{1-a}{b} = r^*.$$

($\Leftarrow$) Conversely, assume $r = r^* = (1-a)/b$. Then

$$a + rb = a + \frac{1-a}{b} b = a + (1-a) = 1.$$

By part (i), $Z^{\text{pooled}}/Z_A = a + rb = 1$. Multiplying by $Z_A > 0$ yields $Z^{\text{pooled}} = Z_A$. This completes the iff proof.

Step 3: Proof of (iii) (Sign characterization and indeterminacy). Because $Z^{\text{pooled}} > 0$ and $Z_A > 0$, the ratio Z^{pooled}/Z_A is positive and ln is defined. Moreover, $\ln(\cdot)$ is strictly increasing on $(0, \infty)$, so for any $x > 0$,

$$\ln(x) \gtrless 0 \quad \iff \quad x \gtrless 1.$$

Apply this with $x = Z^{\text{pooled}}/Z_A$:

$$\ln Z^{\text{pooled}} - \ln Z_A = \ln\left(\frac{Z^{\text{pooled}}}{Z_A}\right) \gtrless 0 \quad \iff \quad \frac{Z^{\text{pooled}}}{Z_A} \gtrless 1.$$

Since $Z_A > 0$, multiplying the inequality $\frac{Z^{\text{pooled}}}{Z_A} \gtrless 1$ by Z_A preserves the direction, giving

$$\frac{Z^{\text{pooled}}}{Z_A} \gtrless 1 \quad \iff \quad Z^{\text{pooled}} \gtrless Z_A.$$

Finally, by (i) we have $Z^{\text{pooled}}/Z_A = a + rb$ with $b > 0$. Thus

$$\frac{Z^{\text{pooled}}}{Z_A} \gtrless 1 \quad \iff \quad a + rb \gtrless 1 \quad \iff \quad rb \gtrless 1 - a \quad \iff \quad r \gtrless \frac{1-a}{b} = r^*,$$

where the last equivalence uses $b > 0$.

For the indeterminacy statement: for fixed $(\beta_k, \beta_l, \beta_m, s_K, s_L, s_M)$, the threshold r^* is fixed and positive. Since r can take values in $(0,\infty)$, there exist values $r<r^*$, $r=r^*$, and $r>r^*$. By the proven equivalences, consolidation bias can therefore be negative, zero, or positive without additional restrictions on r. ■

G.2 Corollary 1 (CRS direction)

Corollary 1 (Direction of consolidation bias under constant returns to scale)**.** *Maintain the assumptions and notation of Proposition 1, and suppose $\beta_k+\beta_l+\beta_m=1$ (constant returns to scale). Then*

$$r^* \;\geq\; 1,$$

with equality if and only if $s_K=s_L=s_M$. Consequently, the consolidation bias is weakly downward ($Z^{\text{pooled}}\leq Z_A$) whenever the acquirer is at least as productive as the target ($r\leq 1$), and remains weakly downward for all $r\leq r^*$.

Proof. Define $g(x_1,x_2,x_3)=x_1^{\beta_k}\,x_2^{\beta_l}\,x_3^{\beta_m}$. Under constant returns to scale, g is concave and homogeneous of degree one on $\mathbb{R}^3_{++}$, and therefore superadditive: $g(\mathbf{x}+\mathbf{y})\geq g(\mathbf{x})+g(\mathbf{y})$ for all $\mathbf{x},\mathbf{y}>0$.[10] Applying this with $\mathbf{x}=(s_K,s_L,s_M)$ and $\mathbf{y}=(1-s_K,1-s_L,1-s_M)$,

$$1=g(1,1,1)=g(\mathbf{x}+\mathbf{y})\;\geq\;g(\mathbf{x})+g(\mathbf{y})=a+b.$$

Since $a+b\leq 1$, it follows that $1-a\geq b>0$ and hence $r^*=(1-a)/b\geq 1$.

Equality $a+b=1$ holds if and only if the concavity step binds, which requires $\mathbf{x}/\|\mathbf{x}\|_1=\mathbf{y}/\|\mathbf{y}\|_1$, i.e., $\mathbf{x}$ and $\mathbf{y}$ are proportional. Since $x_j=s_j$ and $y_j=1-s_j$, proportionality requires $s_j/(1-s_j)$ to be constant across $j\in\{K,L,M\}$, which holds if and only if $s_K=s_L=s_M$.

The final claim follows from Proposition 1(iii): $r\leq 1\leq r^*$ implies $r\leq r^*$ and hence $\ln Z^{\text{pooled}}-\ln Z_A\leq 0$. ■

G.3 Corollary 2 (Special cases)

Corollary 2 (Special cases of Proposition 1)**.** *Maintain the assumptions and notation of Proposition 1.*

(i) ***Constant input ratios.*** *Suppose $s_K=s_L=s_M\equiv s\in(0,1)$, or equivalently $\lambda\equiv K_A/K_T=L_A/L_T=M_A/M_T>0$, where $s=\lambda/(\lambda+1)$. Let $\beta\equiv\beta_k+\beta_l+\beta_m>0$. Then $a=s^\beta$, $b=(1-s)^\beta$, and*

$$Z^{\text{pooled}}=Z_A\cdot\frac{\lambda^\beta+r}{(\lambda+1)^\beta},$$

so the consolidation bias is characterized by the threshold

$$Z^{\text{pooled}}\gtrless Z_A\quad\Longleftrightarrow\quad r\gtrless(\lambda+1)^\beta-\lambda^\beta.$$

[10]For any concave g homogeneous of degree one, write $g(\mathbf{x}+\mathbf{y})=\|\mathbf{x}+\mathbf{y}\|_1\,g\big(\frac{\mathbf{x}+\mathbf{y}}{\|\mathbf{x}+\mathbf{y}\|_1}\big)\geq\|\mathbf{x}+\mathbf{y}\|_1\big[\frac{\|\mathbf{x}\|_1}{\|\mathbf{x}+\mathbf{y}\|_1}\,g\big(\frac{\mathbf{x}}{\|\mathbf{x}\|_1}\big)+\frac{\|\mathbf{y}\|_1}{\|\mathbf{x}+\mathbf{y}\|_1}\,g\big(\frac{\mathbf{y}}{\|\mathbf{y}\|_1}\big)\big]=g(\mathbf{x})+g(\mathbf{y})$, where the inequality is concavity and the equalities use homogeneity of degree one.

Under constant returns to scale ($\beta = 1$), the threshold equals one and $\text{sign}(Z^{\text{pooled}} - Z_A) = \text{sign}(Z_T - Z_A)$. *Under decreasing returns to scale ($\beta < 1$), the threshold is strictly below one, so the consolidation bias is upward at productivity parity ($r = 1$) and can be upward even when the acquirer is more productive than the target. This contrasts with the CRS result in Corollary 1.*

(ii) ***Symmetric firms under constant returns.*** *Suppose $\beta = 1$ and the acquirer and target are of equal size ($\lambda = 1$, equivalently $s = 1/2$). Then $a = b = 1/2$ and*

$$Z^{\text{pooled}} = \frac{Z_A + Z_T}{2}.$$

Pooled TFPR is the arithmetic average of the two firms' individual TFPRs. The direction of the consolidation bias is determined entirely by the acquirer–target productivity gap:

(a) *if $Z_A > Z_T$, then $Z^{\text{pooled}} < Z_A$ and the bias is downward;*

(b) *if $Z_A = Z_T$, then $Z^{\text{pooled}} = Z_A$ and the bias is zero;*

(c) *if $Z_A < Z_T$, then $Z^{\text{pooled}} > Z_A$ and the bias is upward.*

(iii) ***Negligible target.*** *As the target's share of all inputs vanishes ($s_K, s_L, s_M \to 1$), $a \to 1$ and $b \to 0$, so*

$$\frac{Z^{\text{pooled}}}{Z_A} = a + r\,b \;\to\; 1$$

regardless of the target's relative productivity r. The consolidation bias vanishes when the target is small relative to the acquirer.

Proof. *Part (i).* Assume $s_K = s_L = s_M \equiv s$, equivalently $\lambda \equiv K_A/K_T = L_A/L_T = M_A/M_T > 0$, where $s = \lambda/(\lambda + 1)$. Then $K_A = \lambda K_T$, $L_A = \lambda L_T$, and $M_A = \lambda M_T$, so under Cobb–Douglas,

$$Q_A = Z_A \lambda^\beta K_T^{\beta_k} L_T^{\beta_l} M_T^{\beta_m}, \qquad Q_T = Z_T K_T^{\beta_k} L_T^{\beta_l} M_T^{\beta_m},$$

where $\beta \equiv \beta_k + \beta_l + \beta_m$. Pooled output and inputs satisfy

$$Q^{\text{pooled}} = (Z_A\lambda^\beta + Z_T) K_T^{\beta_k} L_T^{\beta_l} M_T^{\beta_m}, \qquad (K^{\text{pooled}})^{\beta_k} (L^{\text{pooled}})^{\beta_l} (M^{\text{pooled}})^{\beta_m} = (\lambda+1)^\beta K_T^{\beta_k} L_T^{\beta_l} M_T^{\beta_m}.$$

Cancelling $K_T^{\beta_k} L_T^{\beta_l} M_T^{\beta_m} > 0$ gives

$$Z^{\text{pooled}} = Z_A \cdot \frac{\lambda^\beta + r}{(\lambda + 1)^\beta}, \qquad r \equiv \frac{Z_T}{Z_A}.$$

Since $(\lambda + 1)^\beta > 0$ and $Z_A > 0$,

$$Z^{\text{pooled}} \gtreqless Z_A \quad \iff \quad \lambda^\beta + r \gtreqless (\lambda + 1)^\beta \quad \iff \quad r \gtreqless (\lambda + 1)^\beta - \lambda^\beta.$$

For the CRS sub-case ($\beta = 1$): the threshold equals $(\lambda+1) - \lambda = 1$, and

$$Z^{\text{pooled}} - Z_A = \frac{\lambda Z_A + Z_T - (\lambda+1)Z_A}{\lambda+1} = \frac{Z_T - Z_A}{\lambda+1},$$

so $\text{sign}(Z^{\text{pooled}} - Z_A) = \text{sign}(Z_T - Z_A)$ since $\lambda + 1 > 0$.

For the DRS sub-case ($\beta < 1$): define $h(\lambda) = (\lambda+1)^\beta - \lambda^\beta$ for $\lambda > 0$. Then $h(0) = 1$ and $h'(\lambda) = \beta[(\lambda+1)^{\beta-1} - \lambda^{\beta-1}] < 0$ because $\beta - 1 < 0$ makes $x^{\beta-1}$ strictly decreasing, so $(\lambda+1)^{\beta-1} < \lambda^{\beta-1}$. Since h is strictly decreasing from $h(0) = 1$, the threshold $(\lambda+1)^\beta - \lambda^\beta < 1$ for all $\lambda > 0$. At productivity parity ($r = 1$), we therefore have $r = 1 > (\lambda+1)^\beta - \lambda^\beta$, so $Z^{\text{pooled}} > Z_A$ and the bias is upward.

Part (ii). Set $\beta = 1$ and $\lambda = 1$ (equivalently $s = 1/2$). From part (i),

$$Z^{\text{pooled}} = Z_A \cdot \frac{1+r}{2} = \frac{Z_A + Z_T}{2}.$$

The three cases follow immediately: (a) if $Z_A > Z_T$, then $(Z_A + Z_T)/2 < Z_A$; (b) if $Z_A = Z_T$, then $(Z_A + Z_T)/2 = Z_A$; (c) if $Z_A < Z_T$, then $(Z_A + Z_T)/2 > Z_A$.

Part (iii). From Proposition 1(i), $Z^{\text{pooled}}/Z_A = a + rb$ with $a = s_K^{\beta_k} s_L^{\beta_l} s_M^{\beta_m}$ and $b = (1-s_K)^{\beta_k}(1-s_L)^{\beta_l}(1-s_M)^{\beta_m}$. As $s_K, s_L, s_M \to 1$, continuity of x^{β_j} gives $a \to 1$ and $b \to 0$. Hence $a + rb \to 1$ for any fixed $r > 0$, and $Z^{\text{pooled}}/Z_A \to 1$. ■

G.4 Proposition 2 (Markup aggregation)

We restate the proposition here with all notation defined.

Proposition 2 *(Sales-weighted aggregation exceeds the pooled markup).* Fix a deal and year t. Suppose $Q_{A,t} > 0$, $Q_{T,t} > 0$, and material expenditures satisfy $M_{A,t} > 0$ and $M_{T,t} > 0$. Assume a common output elasticity of materials $\beta_m > 0$ for the acquirer and target (as in horizontal deals). Define firm-level markups and material expenditure shares by

$$\mu_{i,t} = \frac{\beta_m}{c_{i,t}}, \qquad c_{i,t} \equiv \frac{M_{i,t}}{Q_{i,t}}, \qquad i \in \{A, T\},$$

sales weights by

$$\pi_{A,t} \equiv \frac{Q_{A,t}}{Q_{A,t} + Q_{T,t}}, \qquad \pi_{T,t} \equiv 1 - \pi_{A,t},$$

the sales-weighted markup by

$$\mu_t^{\text{sales}} \equiv \pi_{A,t}\,\mu_{A,t} + \pi_{T,t}\,\mu_{T,t},$$

and the pooled (consolidated-accounts) markup by

$$\mu_t^{\text{pooled}} \equiv \beta_m \cdot \frac{Q_{A,t} + Q_{T,t}}{M_{A,t} + M_{T,t}} = \frac{\beta_m}{\pi_{A,t}\,c_{A,t} + \pi_{T,t}\,c_{T,t}}.$$

Then

$$\mu_t^{\text{sales}} \;\geq\; \mu_t^{\text{pooled}},$$

with equality if and only if $c_{A,t} = c_{T,t}$ (equivalently, the materials expenditure shares are identical across the acquirer and target).

Proof. Let $c_{i,t} \equiv \frac{M_{i,t}}{Q_{i,t}}$ and define $f(c) \equiv \beta_m / c$ for $c > 0$. Then

$$f''(c) = \frac{2\beta_m}{c^3} > 0,$$

so f is strictly convex on $(0, \infty)$. Let sales weights be

$$\pi_{A,t} \equiv \frac{Q_{A,t}}{Q_{A,t} + Q_{T,t}}, \qquad \pi_{T,t} \equiv 1 - \pi_{A,t}.$$

Assuming $Q_{A,t} > 0$ and $Q_{T,t} > 0$, we have $\pi_{A,t}, \pi_{T,t} \in (0,1)$. Using $\mu_{i,t} = f(c_{i,t})$ and Jensen's inequality,

$$\mu_t^{\text{sales}} = \pi_{A,t}\mu_{A,t} + \pi_{T,t}\mu_{T,t} = \pi_{A,t} f(c_{A,t}) + \pi_{T,t} f(c_{T,t}) \;\geq\; f(\pi_{A,t}c_{A,t} + \pi_{T,t}c_{T,t}).$$

Moreover,

$$\begin{aligned}
\pi_{A,t}c_{A,t} + \pi_{T,t}c_{T,t} &= \frac{Q_{A,t}}{Q_{A,t} + Q_{T,t}} \frac{M_{A,t}}{Q_{A,t}} + \frac{Q_{T,t}}{Q_{A,t} + Q_{T,t}} \frac{M_{T,t}}{Q_{T,t}} \\
&= \frac{M_{A,t} + M_{T,t}}{Q_{A,t} + Q_{T,t}}.
\end{aligned}$$

Therefore,

$$f(\pi_{A,t}c_{A,t} + \pi_{T,t}c_{T,t}) = \frac{\beta_m}{\frac{M_{A,t}+M_{T,t}}{Q_{A,t}+Q_{T,t}}} = \beta_m \cdot \frac{Q_{A,t} + Q_{T,t}}{M_{A,t} + M_{T,t}} = \mu_t^{\text{pooled}}.$$

Combining these displays yields $\mu_t^{\text{sales}} \geq \mu_t^{\text{pooled}}$. Finally, since f is strictly convex and $\pi_{A,t}, \pi_{T,t} \in (0,1)$, equality in Jensen's inequality holds if and only if $c_{A,t} = c_{T,t}$, i.e.,

$$\frac{M_{A,t}}{Q_{A,t}} = \frac{M_{T,t}}{Q_{T,t}}.$$

■