Maastricht University

UNU MERIT

Working Paper Series

#2026-017

Superstars Have the Power? Technological M&As and Firm Markups

Cecilia Seri, Elena Zaurino, Maria Martinez Cillero, Lorenzo Napolitano and Francesco Rentocchini

Published 20 November 2026

DOI: https://www.doi.org/10.53330/NJAB4627

Maastricht Economic and social Research institute on Innovation and Technology (UNU-MERIT)
email: **info@merit.unu.edu** | website: **http://www.merit.unu.edu**

Boschstraat 24, 6211 AX Maastricht, The Netherlands
Tel: (31) (43) 388 44 00

UNU-MERIT Working Papers
ISSN 1871-9872

Maastricht Economic and social Research Institute on Innovation and Technology
UNU-MERIT

UNU-MERIT Working Papers intend to disseminate preliminary results of research carried out at UNU-MERIT to stimulate discussion on the issues raised.



Superstars Have the Power? Technological M&As and Firm Markups

Cecilia Seri[1], Elena Zaurino[2], Maria Martinez Cillero[2], Lorenzo Napolitano[3], and Francesco Rentocchini[4]

[1]*UNU-MERIT and Maastricht University, Maastricht, The Netherlands*
[2]*European Commission, Joint Research Centre (JRC), Ispra, Italy*
[3]*Independent Researcher, Pisa, Italy*
[4]*Department of Economics and Management, University of Ferrara, Italy*

Abstract

We study whether acquisitions of patent-holding firms (technological M&As) raise acquirer markups, exploiting the staggered timing of first technological acquisitions in a panel of 7,796 listed firms in 79 countries over 2009–2020. Technological M&As raise acquirer markups by 1.7 percent, roughly two years of growth in the aggregate US markup, with larger effects among top R&D investors. Post-acquisition patenting shows no detectable increase, and productivity declines. Markup gains increase with the size of the acquired patent stock and with technological proximity to the target, consistent with acquired patents shielding acquirers from rivals rather than with innovation or efficiency gains.

JEL classification: D43, G34, L13, L40, O31, O34

Keywords: mergers and acquisitions; market power; markups; innovation; patents; knowledge diffusion.

Rising markups and the growing dominance of leading firms have raised questions about the sources of increasing market power and its dynamic effects (Autor et al., 2020; De Loecker et al., 2020). Innovation and technological leadership play a key role in these patterns: higher markups may result from productivity improvements and continued innovation by technological leaders, but may also reflect their ability to concentrate knowledge, limit its diffusion, and protect themselves from competition, which weakens incentives to innovate. Acquisitions of innovative firms (technological M&As) can contribute to both channels. On the one hand, they can favor the integration of complementary knowledge bases and the productive deployment of acquired intangibles, which can raise productivity, lower marginal costs, and stimulate further invention (Bena and Li, 2014). On the other hand, they can be used strategically to accumulate patent portfolios restricting access to knowledge domains and insulating the acquirer from rivals (Hall and Ziedonis, 2001; Shapiro, 2000). Both channels can raise acquirer markups, though with different welfare implications. Despite their relevance for competition, innovation, and aggregate markups, systematic evidence on the effect of technological M&As on acquirer markups, and on the channels behind it, is scarce and confined to specific settings. Estimates across sectors and countries are needed to relate these acquisitions to aggregate markup trends and to provide firm-level evidence on the channels linking innovation and market power (Akcigit and Ates, 2023; De Ridder, 2024).

Against this backdrop, we ask whether technological M&As raise acquirer markups and examine the channels underlying this effect, spanning more industries and countries than earlier studies. To that end, we assemble a panel of 7,796 publicly listed firms in 79 countries over 2009–2020. We link acquisition records, ownership structures, financial statements, and patent histories from Orbis to follow acquirers and the patent portfolios of their subsidiaries. We define a technological M&A as the acquisition of a target holding at least one patent family. Our main outcome is the acquirer's markup, the ratio of price to marginal cost, estimated using the production-function approach of De Loecker and Warzynski (2012).

We leverage variation in the timing of firms' first technological acquisition to estimate the causal effect of technological M&As on acquirer markups. Our difference-in-differences design compares the markup trajectories of these firms with those of acquirers that have not yet completed a technological M&A or never do so during the sample period. Restricting the comparison to acquirers limits pre-existing differences associated with acquisition activity (David, 2021). After accounting for persistent differences between firms and common industry-specific shocks, we find no evidence of differential markup trends before the deal.

Our results show that technological M&As raise acquirer markups by 1.7 percent on average. The effect persists over the observed post-acquisition period and is largest

at the longest horizons. Evaluated at the pre-treatment mean, the estimate raises the margin over marginal cost from about 27.8 to 29.1 cents per unit of revenue. This is equivalent to roughly two years of growth in the aggregate US markup documented by De Loecker et al. (2020). The results are robust to estimators that allow treatment effects to vary across cohorts (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021) and to restricting the comparison to later technological acquirers, whose observed pre-acquisition characteristics are close to balanced. They also hold in specifications controlling for pre-acquisition capital stock, firm age, patent stock, and operating revenue, and when matching on industry and acquisition year (technological acquisition vs general acquisition). Operating profits also rise, and markup gains persist after adjustment for an industry-level fixed-cost benchmark, supporting a market-power interpretation of rising markups.

What accounts for these higher markups? The productive deployment of acquired knowledge could make acquirers more inventive or more efficient (the innovation-enhancing channel). Yet we find no statistically detectable increase in patenting by the acquirer and its subsidiaries, while productivity declines in the years immediately after the acquisition. We also find limited evidence of a shift toward fixed costs associated with intangible-driven economies of scale (De Ridder, 2024). These findings provide little support for innovation enhancement or efficiency gains as explanations for the markup increase. Markup gains that do not arise from the use of acquired knowledge may instead arise from its ownership, as acquired patents can protect the acquirer from competition (the innovation-shielding channel). We therefore relate markups to the size of the acquired patent portfolio, since larger portfolios can more strongly restrict the freedom of rivals to operate, further strengthening the bargaining position of the acquirer (Czarnitzki et al., 2020; Hall and Ziedonis, 2001). We also relate markups to the proximity between acquirer and target, given that overlap in technology and product markets can concentrate patent rights around the core technologies of the acquirer and remove an actual or potential rival (Cunningham et al., 2021). Markups are higher where the acquired portfolio is larger, and the estimated markup increase is more than twice as large among acquirers with high technological and product-market proximity to their targets as among other technological acquirers. Larger gains also occur when technological proximity is high but product-market proximity is low. Read together, these patterns are consistent with the consolidation of patent rights around the core technologies of the acquirer strengthening its position vis-à-vis rivals (Hall and Ziedonis, 2001; Shapiro, 2000).

The effect is larger among top R&D investors (2.6 percent) than among other firms (0.8 percent), a statistically significant difference ($p = 0.010$), which is in line with the concentration of market power among leading firms (Autor et al., 2020; De Loecker et al., 2020). Geographically, the effect is concentrated among firms headquartered

in the United States (3.7 percent) and the European Union (3.5 percent). Sectorally, it is concentrated in high- and medium-high-tech manufacturing (2.7 and 1.3 percent, respectively) and knowledge-intensive services (2.9 percent), where proprietary technology and intangible assets are a main source of competitive advantage (Crouzet et al., 2022). The point estimates are also larger for acquisitions involving digital and AI patents, at 3.4 and 5.1 percent, respectively. The AI estimate is suggestive, however, as it rests on few treated firms and shows evidence of differential pre-acquisition trends.

Taken together, the evidence is more consistent with technological M&As insulating acquirers from rivals than with innovation or efficiency gains. It does not, however, identify the motives of acquirers or distinguish between the effects of a denser patent portfolio and the removal of an actual or potential rival.

This paper contributes to three strands of literature. First, it relates to work on how mergers and acquisitions affect market power and innovation. While technological change is a leading explanation for the rise in market power (Akcigit and Ates, 2023; Autor et al., 2020; De Loecker and Eeckhout, 2026), studies of mergers have often treated the two outcomes as separate questions. Studies of market power measure merger effects on prices and markups in single industries, where product-level prices and models of demand and competition allow precise measurement (Bhattacharya et al., 2026; Miller and Weinberg, 2017), and for the targets or rivals of merging firms (Stiebale and Szücs, 2022). Other studies measure the contribution of mergers to the rise in aggregate markups (Cao and Zhu, 2025; Levonyan and Mengano, 2024). Studies of innovation show that acquisitions can discontinue overlapping projects and reduce patenting (Cunningham et al., 2021; Haucap et al., 2019; Seru, 2014), deter entry (Kamepalli et al., 2020), or instead encourage small firms to innovate in order to be acquired (Phillips and Zhdanov, 2013), while their effects on the innovation of acquirers are mixed (Bena and Li, 2014; Szücs, 2014). The few studies that measure both outcomes focus on the rivals of merging firms (Stiebale and Szücs, 2022), and the evidence more generally comes from single industries, from a few large acquirers (Ederer et al., 2026; Jin et al., 2023), or from mergers reviewed by competition authorities (Stiebale and Szücs, 2026). Potentially anti-competitive acquisitions, however, can fall below mandatory notification thresholds and consequently escape regulatory scrutiny (Wollmann, 2019). We complement these studies by estimating, for the same acquirers and across industries and countries, how the acquisition of patent-filing targets affects both markups and innovation. Our production-function markups are less precise than product-level measures, but they can be estimated consistently across this range of settings and are widely used to measure market power across firms and countries (De Loecker et al., 2020; De Ridder et al., 2026). Because our deal data record majority acquisitions irrespective of their size, the sample is not limited to transactions notified to competition authorities.

Second, the paper relates to work on the rise of market power and its sources. The aggregate increase in markups reflects a reallocation of activity toward high-markup firms and the growth of superstar firms (Autor et al., 2020; De Loecker et al., 2020). Proposed explanations emphasize intangible and proprietary technology assets (Bessen, 2020; De Ridder, 2024) and slower diffusion of knowledge from frontier firms to laggards (Akcigit and Ates, 2021, 2023). Much of this evidence is aggregate or industry-level, or else structural or correlational at the firm level, and causal evidence on the specific actions through which firms raise their markups is scarce. We provide such evidence for a specific corporate strategy, the acquisition of patent-filing firms, which raises acquirer markups, more so among top R&D investors. The association of these gains with the acquired patent stock and with technological proximity, together with the absence of a shift toward fixed costs, is more consistent with explanations based on slower knowledge diffusion than with intangible-driven economies of scale.

Third, the paper relates to work on the strategic use of patents. Patents confer exclusion rights that firms can use beyond protecting their own inventions: to build portfolios that strengthen their bargaining position (Hall and Ziedonis, 2001; Shapiro, 2000) and to block rivals (Czarnitzki et al., 2020). Empirical studies of strategic patenting focus mostly on the patents firms file themselves: in biopharmaceuticals, dense webs of later-expiring secondary patents delay the entry of competitors (Frakes and Wasserman, 2025), and across industries, strategic patents are associated with higher market concentration and with lower productivity and patenting of competitors (Veihl, 2022). Patents can also be acquired: firms with distinctive technology portfolios become prime targets for close product-market rivals seeking to reduce competition and for technologically proximate acquirers (Arts et al., 2025), but the consequences of such acquisitions for the market power of the buyer are less documented. We show that markup gains are larger where the acquired portfolio is larger and technologically closer to that of the acquirer, linking the market for patent-filing firms to the market power of buyers. This evidence is consistent with the strategic accumulation of patents, although it does not identify the motives of acquirers or separate the effect of a denser patent portfolio from that of removing a rival.

The remainder of the paper is organized as follows. Section I describes the data and the estimation sample, Section II presents the empirical strategy and the mechanism tests, Section III reports the baseline, heterogeneity, and mechanism results, Section IV assesses the economic magnitude of the effects, and Section V concludes. The Supplemental Appendix describes the markup estimation and the construction of the patent data, and reports additional descriptive statistics and robustness checks.

I Data

Assessing how technological M&As affect the market power of acquiring firms requires linking acquisition records to firms' pre- and post-acquisition financial information and to the patent stocks and flows of acquirers and their subsidiaries, including newly acquired firms. To build this dataset, we mainly rely on Orbis, maintained by Moody's (formerly Bureau van Dijk), combining information on M&As, ownership structures, firm accounts, and patents through a common firm identifier. The estimation sample covers 7,796 listed firms in 79 countries that conduct at least one M&A during 2009–2020.

A. Data for M&A and Firm Financial Information

We combine M&A records from Zephyr[1] with ownership information from Orbis Historical Ownership and financial statements from Orbis Global. We identify majority (i.e., control) acquisitions over 2007–2020 using information on the initial, acquired, and final equity stakes, retaining transactions in which the acquirer's final stake, or the stake acquired in the transaction, exceeds 50 percent.[2]

For each acquirer, we reconstruct its ownership structure annually. We first identify its subsidiaries in 2007, the earliest year available in Orbis Historical Ownership, retaining firms for which the acquirer is the global ultimate owner or direct majority shareholder, using a 50 percent ownership threshold. We then update this initial structure using acquisitions completed between 2008 and 2020, tracking the transfer of subsidiaries across acquirers.[3] This procedure identifies the subsidiaries belonging to each acquirer in each year, accounting for subsidiary transfers through acquisitions.

We then extract acquirers' financial statements from Orbis Global to estimate firm-level markups, our measure of market power and main outcome variable of interest.[4]

[1] Zephyr is the M&A database associated with Orbis and draws on media reports, press releases, financial reports, and administrative filings.

[2] We exclude deals that lack stake information and are classified as minority acquisitions below 10 percent, deals in which both the initial and the final stake are reported as 100 percent (reported stakes are rounded, so these deals show no effective increase), and deals in which the acquirer already held a controlling stake above 50 percent before the transaction. We retain transactions in which either the final stake or the stake acquired in the transaction exceeds 50 percent. Using the reported deal type, we also exclude initial public offerings, institutional buyouts, and mergers: respectively, 4 (less than 0.01 percent), 15,459 (4.93 percent), and 2,121 (0.68 percent) deals. Intra-company deals and joint ventures are excluded from the sample.

[3] When an acquirer is itself acquired, we assign its subsidiaries to the new parent, representing ownership at a single level.

[4] Limited financial data prevent markup estimation in the following NACE Rev. 2 sectors, which we therefore exclude: agriculture, forestry and fishing (01–03); mining and quarrying (05–09); postal and courier activities (53); residential care (87); repair of computers and personal and household goods (95); household activities (97–98); and activities of extraterritorial organizations and bodies (99). Markup studies based on European balance-sheet data apply similar restrictions, confining the sample to the non-

We retain one observation per firm-year, prioritizing unconsolidated accounts over consolidated accounts and excluding accounts with limited financial information. Markups, the ratio of price to marginal cost, capture firms' ability to price above marginal cost. Unlike concentration ratios, their estimation using the production-function approach does not require defining product-market boundaries, making them useful for our cross-country, cross-industry analysis (Syverson, 2019). Following De Loecker and Warzynski (2012), we estimate this ratio as the output elasticity of intermediate inputs divided by their expenditure share in firm revenue, treating intermediate inputs as flexible. Because our design exploits within-firm changes in markups, biases in markup levels from revenue-based estimation are of less concern (De Ridder et al., 2026). We estimate translog production functions separately by two-digit NACE sector using the control-function approach of Ackerberg et al. (2015) to account for the endogeneity of input choices. We use revenue to proxy for output, total fixed assets (including tangible and intangible assets) for capital, employee costs for labor, and material expenses for intermediate inputs.[5] Monetary variables are deflated using country-level GDP deflators from the World Bank. Supplemental Appendix A1 presents the theoretical framework, while Supplemental Appendix A2 describes the markup estimation, account-selection criteria, and data-cleaning procedures, with detailed variable definitions in Supplemental Appendix Table A1. Supplemental Appendix D1 examines the sensitivity of our results to changes in account consolidation status: the estimate is unchanged when controlling for consolidation status, and remains positive and significant among firms that do not switch between consolidated and unconsolidated accounts.

We complement the markup analysis with operating profitability and an excess-markup measure that adjusts for an industry-by-year benchmark fixed-cost share. These checks help assess whether the estimated markup increases are accompanied by higher profitability and persist after accounting for fixed-cost coverage (Supplemental Appendix D2).

B. Patent Data Sources

We use patent data on acquirers and their subsidiaries to identify technological M&As and examine post-acquisition patenting activity and the role of acquired patent portfolios.

We rely on Orbis Intellectual Property (IP), which provides worldwide information

financial business economy and excluding agriculture, mining, and non-market services (e.g., Calligaris et al., 2024).

[5]When material expenses are missing but cost of goods sold is available, we impute them by multiplying the firm's cost of goods sold by the average materials-to-cost ratio among observations reporting both variables in the same two-digit NACE sector, pooling across years.

on patent applications.[6] We use applications from 1970 to 2020 to reconstruct patent histories before the acquisition events and count patent families to avoid double-counting the same invention across patent offices. Combining these records with the annual ownership structures described above, we construct: (i) the annual flow of patent applications by each acquirer and its subsidiaries; (ii) the patent stock held by each acquirer and its subsidiaries in each year, including portfolios obtained through acquisitions; and (iii) the stock of patent applications filed by each target up to the year of acquisition. We classify an acquisition as a technological M&A if the target holds at least one patent family in its stock at the time of the deal.

We also construct indicators of international patent protection and digital technologies, and a measure of technological proximity between acquirers and targets, to examine heterogeneity in the markup effects of technological M&As.

First, we use Orbis IP to identify IP5 patent families, defined as families with filings in at least two patent offices worldwide, including at least one of the five largest offices: the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the United States Patent and Trademark Office (USPTO), or the China National Intellectual Property Administration (CNIPA). We use IP5 status as a proxy for patent value, since seeking protection across multiple patent offices entails additional costs that applicants are more likely to incur for commercially valuable inventions. Details on patent-office identification and country-code harmonization are provided in Supplemental Appendix B1.

Second, we identify patents related to digital general-purpose technologies (dGPTs) and artificial intelligence (AI) using patent-level tags provided by WIPO,[7] based on the classification developed in WIPO (2022) and updated in 2023. This classification combines International Patent Classification (IPC) and Cooperative Patent Classification (CPC) codes with keyword searches. We link these tags to Orbis IP and classify a patent family as AI if it contains at least one AI-tagged patent, and as digital if it contains at least one patent tagged in the remaining dGPT categories. These indicators allow us to examine whether the markup effects of technological M&As differ for acquisitions involving digital and AI technologies. Details on the categories and matching procedure are provided in Supplemental Appendix B2.

Finally, we measure technological proximity between acquirers and targets using the cosine similarity of binary vectors indicating the IPC classes represented in their patent portfolios at the time of the deal. This measure allows us to examine whether markup gains depend on the degree of technological overlap between the merging firms. Our baseline measure uses four-digit IPC codes; we also consider alternative aggregation

[6] We retain patents whose applicants have a firm identifier, allowing us to link patent records to acquisition and financial data.

[7] The World Intellectual Property Organization (WIPO) is the source of the patent-level classification of digital and AI technologies used in this paper; we thank WIPO for sharing these data.

levels and Jaccard similarity as robustness checks. Details on the construction of these measures are provided in Supplemental Appendix B3. We measure product-market proximity analogously, as the cosine similarity of binary vectors indicating the four-digit NACE sectors of the acquirer and the target.

C. Estimation Sample and Descriptive Statistics

We merge the acquisition, ownership, and financial dataset described in Section IA with the patent data using the common firm identifier to construct the estimation sample. We restrict the sample to firms observed for at least six years in the 2008–2020 panel and exclude firms whose first technological M&A occurs within their first four observed years, ensuring an observed pre-treatment history before the first technological acquisition. We also exclude firms first treated in 2020, given the small size of this cohort and the potential disruption associated with the COVID-19 pandemic, and firms with missing financial information in the year of their first technological acquisition. Because the capital stock is constructed from investment, which requires the capital of the previous year, the first observed year of each firm does not enter the estimation sample, which therefore covers 2009–2020.

Our final estimation sample includes 73,690 firm-year observations on 7,796 listed firms in 79 countries over 2009–2020; 1,273 of these firms complete at least one technological M&A between 2012 and 2019.

Two examples illustrate the distinction between treated and comparison firms. Priceline (now Booking Holdings) enters the treated group in 2013 following its acquisition of KAYAK, a patent-owning travel-search company whose technology and user-interface innovations were part of the announced rationale for the deal (KAYAK Software Corporation, 2012). Netflix, by contrast, belongs to the comparison group: although it acquired the comic-book publisher Millarworld in 2017, it has no recorded acquisition of a patent-owning firm during our sample period.[8]

Table 1 compares technological acquirers in the year before their first technological M&A with eligible comparison firms.

Before their first technological acquisition, acquirers have higher operating revenue, are older, and exhibit greater patent ownership and patenting activity than comparison firms, whereas their ratio of capital to total assets is similar. Their average markup is nevertheless lower: 1.38, compared with 1.44 among comparison firms. This gap partly reflects differences in industry composition: reweighting comparison firms to the industry composition of technological acquirers within each acquisition cohort reduces

[8]Netflix's classification is consistent with its development of streaming technology internally alongside acquisitions of content-related intellectual property. It launched its own content-delivery network, Open Connect, in 2012 and described the Millarworld acquisition as a means of developing and owning original content (Florance, 2012; Netflix, 2017).

Table 1: Pre-treatment Characteristics of Technological Acquirers and Comparison Firms

	Technological acquirers	Comparison firms	Normalized difference
Panel A. Levels in the year before the first technological M&A			
Markup	1.379	1.441	-0.21
	(0.277)	(0.312)	
Real operating revenue (millions)	1,577.31	991.44	0.19
	(3,118.91)	(2,982.49)	
	[454.75]	[193.89]	
Capital stock / total assets	0.174	0.170	0.01
	(0.189)	(0.518)	
Firm age (years)	36.45	31.04	0.19
	(30.81)	(25.99)	
Patent stock (PIM)	343.95	59.47	0.25
	(1,516.26)	(600.90)	
	[10.38]	[0.00]	
Patent applications (families per year)	28.02	4.32	0.22
	(142.42)	(55.95)	
	[1.00]	[0.00]	
Firms holding patents (%)	78.31	44.63	0.74
	(41.21)	(49.71)	
Treated firms	1,222		
Comparison observations (firms)		44,989 (7,617)	
Panel B. Changes between three years and one year before the first technological M&A			
Change in markup	-0.012	-0.009	-0.02
	(0.147)	(0.163)	
Revenue growth (log change)	0.118	0.082	0.08
	(0.429)	(0.458)	
Change in patent applications (IHS)	0.088	0.013	0.11
	(0.769)	(0.536)	
Treated firms	1,170		
Comparison observations (firms)		39,176 (7,344)	

Notes: Weighted means, with weighted standard deviations in parentheses and weighted medians in brackets for skewed characteristics. Technological acquirers are observed in the year before their first technological M&A (Panel A) and between three years and one year before it (Panel B). For each acquisition cohort, comparison firms are observed in the same calendar years and comprise acquirers with no recorded technological acquisition during the sample period and firms whose first technological acquisition occurs after the focal cohort year. Treated observations receive weight one; within each cohort and characteristic, comparison weights sum to the number of treated observations with available data. The normalized difference is the difference in means divided by the square root of the average of the two group variances. Of the 1,273 technological acquirers in the estimation sample, 51 have no observation in the year before their first technological M&A. Real operating revenue is winsorized at the 5th and 95th percentiles within industry-year cells. Capital stock is the perpetual-inventory capital stock divided by total assets; patent stock is the perpetual-inventory stock of patent families with 5% annual depreciation. Revenue growth and the change in patent applications are differences in logarithms and in inverse hyperbolic sines. The table describes levels and pre-period changes; it reports neither treatment effects nor tests of parallel trends.

the gap by approximately one-half, although a lower average markup remains.[9] Except for the share of firms holding patents (0.74), all normalized differences in Panel A are below the conventional threshold of 0.25 in absolute value, under which linear regression adjustment is not sensitive to the specification (Imbens and Wooldridge, 2009).

Over the three years preceding the acquisition, changes in markups, revenue, and patenting are similar across the two groups, with normalized differences of at most 0.11 in absolute value. These differences describe the firms selecting into technological acquisitions; their implications for identification are discussed in Section II.

Supplemental Appendix C reports further statistics on the composition of the estimation sample. The number of firms performing their first technological M&A in each year is fairly stable, ranging from 132 to 190 and peaking in 2014 and 2015 (Supplemental Appendix Table C1). The firms in the sample originate from 79 countries representing all global regions; 52.1 percent are located in Asia, followed by Europe (22.1 percent) and North America (19.3 percent) (Supplemental Appendix Table C2). The geographical distribution of the sample closely follows that of all listed firms in Orbis (Supplemental Appendix Figure C1).

II Empirical Strategy

A. Baseline Identification

Does conducting a technological M&A increase firm-level markups? We address this question using a staggered difference-in-differences design that compares the markup trajectories of technological acquirers with those of other acquiring firms. Our baseline event-study specification is:

$$Markup_{it} = \sum_{p=-7}^{-2} \beta_p M\&A_{it}^p + \sum_{p=0}^{5} \beta_p M\&A_{it}^p + \alpha_i + \delta_{s(i),t} + \epsilon_{it} \tag{1}$$

where $Markup_{it}$ is the markup of firm i in year t, with $t = 2009, \ldots, 2020$. Firm fixed effects α_i absorb time-invariant determinants of markups, while industry-by-year fixed effects $\delta_{s(i),t}$ absorb shocks common to firms in the same two-digit NACE industry $s(i)$ in a given year. $M\&A_{it}^p$ are event-time indicators defined relative to firm i's first technological M&A. The indicators for $p = -7$ and $p = 5$ pool observations seven or more years before and five or more years after the acquisition, respectively. We omit

[9] Technological acquirers are more concentrated in manufacturing, particularly electronics, machinery, chemicals, and motor vehicles, where comparison firms have relatively low markups. Comparison firms have greater representation in some higher-markup industries, including real estate. Reweighting reduces the pre-treatment markup gap from approximately 0.062 to 0.027.

the year before the acquisition ($p = -1$), so that each β_p is measured relative to that year. The pre-treatment coefficients provide evidence on differential markup trends before the acquisition. We assess the plausibility of the parallel-trends assumption by examining the pattern of the pre-treatment coefficients $\beta_{-7}, \ldots, \beta_{-2}$ and jointly testing whether they are zero.

The identifying assumption is that, absent a technological M&A, the markups of technological acquirers would have evolved in parallel with those of comparison firms. The main threat is selection: firms that acquire innovative targets may differ from other acquirers in ways that also shape their markup trajectories. In the year before their first technological acquisition, technological acquirers are larger, older, and more patent-intensive than comparison firms, and have lower markups (Table 1). We address this identification threat in multiple ways.

First, the comparison group consists exclusively of acquiring firms. Given that acquirers differ markedly from non-acquirers in size, R&D expenditure, capital investment, and profitability (Bena and Li, 2014; David, 2021; Nindl et al., 2024), firms that undertake non-technological M&A approximate the counterfactual of a technological acquirer more closely than firms that never acquire. Within this population, we consider three nested comparison groups: (i) firms whose first technological M&A occurs later (not-yet-treated firms), together with firms with no technological M&A observed both before and after their own acquisitions; (ii) the same group excluding observations of firms with no technological M&A before their first acquisition, so that comparison firms have already become acquirers and their markup trajectory is not shifted by an acquisition within the estimation window; and (iii) not-yet-treated firms only. The last group compares technological acquirers with firms that complete the same type of acquisition later; in this comparison, pre-treatment characteristics are close to balanced, with normalized differences of at most 0.12 in absolute value (Supplemental Appendix Table C3).

Second, firm fixed effects absorb persistent differences in markup levels, including the lower pre-treatment markup of technological acquirers, and industry-by-year fixed effects absorb shocks common to firms in the same two-digit industry. The latter matter because the markup gap partly reflects the concentration of technological acquirers in manufacturing industries (Section IC).

Third, differences in levels do not by themselves violate the identifying assumption, which concerns changes in markups rather than their levels; they threaten identification to the extent that the characteristics behind them are associated with different markup trajectories (Baker et al., 2026; Kahn-Lang and Lang, 2020). The event-study coefficients for the years before the acquisition test for such differences, and changes in markups, revenue, and patenting over the three years before the acquisition are similar across the two groups (Panel B of Table 1).

Fourth, we assess whether parallel trends hold conditional on the characteristics in which the groups differ. We re-estimate the event study controlling for pre-treatment capital stock, firm age, and patent stock, adding operating revenue in the doubly robust estimator of Sant'Anna and Zhao (2020). We also construct a matched comparison group by Coarsened Exact Matching on two-digit industry and acquisition year (Supplemental Appendix D3).

Fifth, the sample restrictions ensure an observed pre-treatment history. In the 2008–2020 panel, each firm is observed for at least six years, and firms whose first technological M&A occurs within their first four observed years are excluded. Technological acquirers are therefore observed before their first technological acquisition (for at least three years in the estimation sample for 97 percent of them), and their first observed technological acquisition is less likely to follow an earlier, unobserved one.[10]

We exploit variation in the timing of firms' first technological M&A, which occurs between 2012 and 2019 in the estimation sample. We define treatment as absorbing: a firm is treated from the year of its first technological M&A through the end of its observed history. The rationale is that, once a firm has integrated external technologies, patents, and specialized R&D teams, the resultant capabilities and market position are effectively permanent, as displacement of these assets is rarely observed in practice. Moreover, synergies arising from acquired organizational routines and technological know-how can accumulate over time, generating persistent changes in competitive dynamics. Maintaining the treatment indicator at one in all subsequent years accommodates these lasting consequences, while the event-study coefficients allow us to assess whether technological acquisitions increase markups in the deal year and whether this effect persists or strengthens over time. The baseline specification retains firms that undertake multiple technological acquisitions without treating subsequent deals as separate events.[11]

With staggered treatment timing and treatment effects that differ across cohorts, coefficients from a two-way fixed effects (TWFE) event-study specification can reflect treatment effects at other event times (Sun and Abraham, 2021). We therefore report two heterogeneity-robust estimators alongside the TWFE estimates. Following Sun and Abraham (2021), we estimate cohort-specific dynamic effects using the last-treated cohort (2019) and firms with no technological M&A as the comparison group, and

[10] A residual concern is that acquirers may select targets on an unobserved upward technological trajectory, so that part of the post-acquisition markup gain could in principle reflect the organic quality of the target's technology rather than strategic foreclosure. Constructing a comparison group from exogenously withdrawn technological deals would provide a sharper counterfactual against this possibility, but we leave this extension to future work.

[11] In Supplemental Appendix D4, we also estimate dynamic effects using the estimator of De Chaisemartin and d'Haultfoeuille (2024) and an annual indicator equal to one in years in which a firm undertakes a technological M&A and zero otherwise. This specification accommodates repeated changes in acquisition activity.

aggregate them by event time using cohort-share weights. Following Callaway and Sant'Anna (2021), we estimate cohort-by-year average treatment effects on the treated (ATTs) using not-yet-treated firms and firms with no technological M&A as comparison units, and aggregate these effects by event time. The Sun and Abraham specification includes industry-by-year fixed effects; for the Callaway and Sant'Anna estimator, we pre-residualize markups using industry-by-year means estimated among firms with no technological M&A.

B. Mechanism Tests: Outcomes and Predictions

Two distinct mechanisms can explain the increase in acquirer markups following technological M&As. On the one hand, technological M&As can raise markups through the recombination of complementary knowledge bases, R&D synergies, and the productive deployment of acquired intangibles, which can stimulate inventive activity and lower marginal costs (innovation-enhancing channel). On the other hand, a similar increase in markups could result from the consolidation of patent portfolios and the removal of actual or potential rivals, strengthening the acquirer's position in its technological domain (innovation-shielding channel). Both channels are consistent with a positive effect of technological M&As on the acquirer's markup, but have different implications for competition and innovation. No single test can definitively discriminate between them; instead, we examine patenting, productivity, cost structure, acquired patent portfolios, and proximity between merging firms to assess which explanation is more consistent with the evidence.

1. Innovation-Enhancing Channel

To test whether the markup effect of technological M&As operates through innovation enhancement and efficiency gains, we run three complementary tests.

First, we examine the post-acquisition innovation performance of the acquirer using patenting activity as an outcome of technological M&As. We also interact the technological M&A indicator with contemporaneous and subsequent patent flow in the markup regression to examine whether the association between patenting and markups differs in acquisition years.

We measure annual patent flow as the number of patents filed by the merged entity—the acquirer and all its subsidiaries, including the newly acquired firms. When interacting patent flow with the treatment indicator, we consider both patent flow in the deal year and its average over the deal year and the following one, two, or three years. We average over multiple horizons to account for the lag between inventive activity and patent filing, and for the time needed for R&D synergies to materialize and innovation activity to be reorganized following the deal. We consolidate patents across

the merged entity because post-acquisition reorganization of R&D activity and filing entities can make it difficult to attribute patents to individual parts of the combined firm.[12]

If the innovation-enhancing channel operates through R&D synergies, the recombination of complementary knowledge bases should stimulate further inventive activity, showing up as an increase in the acquirer's patenting rate following the deal (Bena and Li, 2014). We therefore apply the baseline difference-in-differences design to annual patent flow, using the absorbing treatment definition based on the first technological acquisition. Following Callaway and Sant'Anna (2021), we estimate cohort-by-year treatment effects, controlling for the acquirer's pre-treatment patent stock, and report their weighted average over the observed post-acquisition years. This estimate measures the average change in annual patenting relative to the estimated counterfactual, expressed in additional patents per year. As a complementary specification, we use a Poisson regression to estimate proportionate changes in expected annual patenting.[13]

Moreover, we estimate a markup regression interacting the technological M&A indicator with patent flow transformed by the inverse hyperbolic sine (IHS). Here, we use a non-absorbing treatment indicator, activated only in each technological acquisition year, to examine the association between patenting and markups in acquisition years.[14] Because patenting may itself respond to the acquisition, we interpret this interaction descriptively rather than as an indirect causal effect of the deal. We estimate:

$$\begin{aligned} Markup_{it} = \beta_1 \, M\&A^{NA}{}_{it} + \beta_2 \, M\&A^{NA}{}_{it} \times PatFlow_{it+h} \\ + \gamma \, PatFlow_{it+h} + \alpha_i + \delta_{s(i),t} + \varepsilon_{it}, \end{aligned} \tag{2}$$

where $M\&A^{NA}{}_{it}$ equals one in each technological acquisition year and $PatFlow_{it+h}$ denotes IHS-transformed average patent flow over year t and the following h years, using available observations ($h \in \{0, 1, 2, 3\}$). For $h = 0$, the measure is patent flow in the deal year alone. Firm fixed effects α_i and industry-by-year fixed effects $\delta_{s(i),t}$ are defined as in equation (1). The coefficient of interest, β_2, measures the difference in the association between patent flow and markups in acquisition years relative to other firm-years.

As a second test of the innovation-enhancing channel, we examine whether technological M&As generate operational efficiency gains for the acquirer by re-estimating

[12]For instance, innovation originating in the acquired subsidiary may be filed under the parent company's name as a result of administrative reassignment rather than any change in underlying inventive activity.

[13]We retain patent counts in the outcome regressions to measure changes in annual patenting on its original scale. In the markup regression, the inverse hyperbolic sine transformation compresses the distribution of patent flow while retaining observations with zero patents.

[14]The baseline markup estimates are similar under the non-absorbing treatment definition using the estimator of De Chaisemartin and d'Haultfoeuille (2024) (Supplemental Appendix Figure D5).

the baseline event-study specification in equation (1) with log total factor productivity (log-TFP) as the outcome variable. We estimate firm-level TFP following the control-function approach of Ackerberg et al. (2015), which addresses the simultaneity bias arising from the correlation between input choices and unobserved productivity shocks. We use materials as the proxy variable and estimate a gross-output production function separately by two-digit NACE sector, recovering log-TFP as the residual from the estimated production function.[15] If the innovation-enhancing channel is operative, the productive deployment of acquired intangibles should translate into a sustained increase in TFP following the deal.

Third, we test whether technological M&As shift the acquirer's cost structure toward fixed costs by applying the baseline difference-in-differences design described in Section IIA to the fixed cost share. Following De Ridder (2024), we proxy the fixed cost share by the difference between the Lerner index ($1 - 1/\mu_{it}$) and the operating margin. Under this cost decomposition, this difference measures fixed costs as a share of revenue. The rationale for this test draws on De Ridder (2024), who shows that intangible inputs lower marginal costs and raise fixed costs, generating economies of scale that can raise markups.[16] An increase in the acquirer's fixed cost share following the deal would be consistent with this form of intangible deployment.

As a complementary check, we also estimate the effect of technological M&As on the acquirer's excess markup to assess whether markup gains persist after adjustment for a fixed-cost benchmark. Excess markup is the Lerner index net of the industry-by-year benchmark fixed-cost share, defined as the leave-one-out, sales-weighted average fixed-cost share in the firm's NACE-by-year cell. This exercise is reported in Supplemental Appendix Table D4.

For TFP and the fixed cost share, we use the absorbing treatment definition of the baseline, on the grounds that integrating acquired technologies can generate persistent changes in productivity and cost structure. For TFP, we replace markups with log-TFP in equation (1) and report TWFE estimates alongside those of Sun and Abraham (2021) and Callaway and Sant'Anna (2021). For the fixed cost share, we estimate average post-acquisition changes using TWFE and Callaway and Sant'Anna (2021), including revenues as a control in selected specifications.

2. Innovation-Shielding Channel

To test whether the markup effect of technological M&As operates through the strategic accumulation of patent portfolios to insulate the acquirer from rivals, we run two

[15] In robustness exercises, we also report results based on the Wooldridge (2009) GMM implementation of ACF, which estimates the production function in a single step and delivers more efficient standard errors. Results are quantitatively and qualitatively in line with those presented in the main text.

[16] Similarly, De Loecker et al. (2020) document rising overhead costs in the United States since the 1980s, but show that these do not fully account for the increase in markups.

complementary tests.

First, we examine whether the stock of patents acquired through technological M&As is associated with higher acquirer markups.

We measure the acquired patent stock as the total number of patents filed by the target up to the year of the deal, which we take as a proxy for the patent portfolio transferred to the acquirer through the acquisition.[17] We accumulate these portfolios across acquisitions and account for obsolescence using the perpetual inventory method with annual depreciation of 5 percent (Hall et al., 2005). The acquired patent stock is IHS-transformed. If the innovation-shielding channel is operative, larger acquired patent portfolios should be associated with higher markups, consistent with the strategic use of patents to limit rivals' freedom to operate and strengthen bargaining positions (Czarnitzki et al., 2020; Hall and Ziedonis, 2001; Shapiro, 2000).

In this test, as in the markup regression with patent flow, we use a non-absorbing treatment indicator, activated only in each technological acquisition year, to account for markup differences associated with new acquisition events. We estimate:

$$Markup_{it} = \beta_1 M\&A^{NA}{}_{it} + \beta_2 PatStock_{it}^{\text{target}} + \alpha_i + \delta_{s(i),t} + \varepsilon_{it}, \tag{3}$$

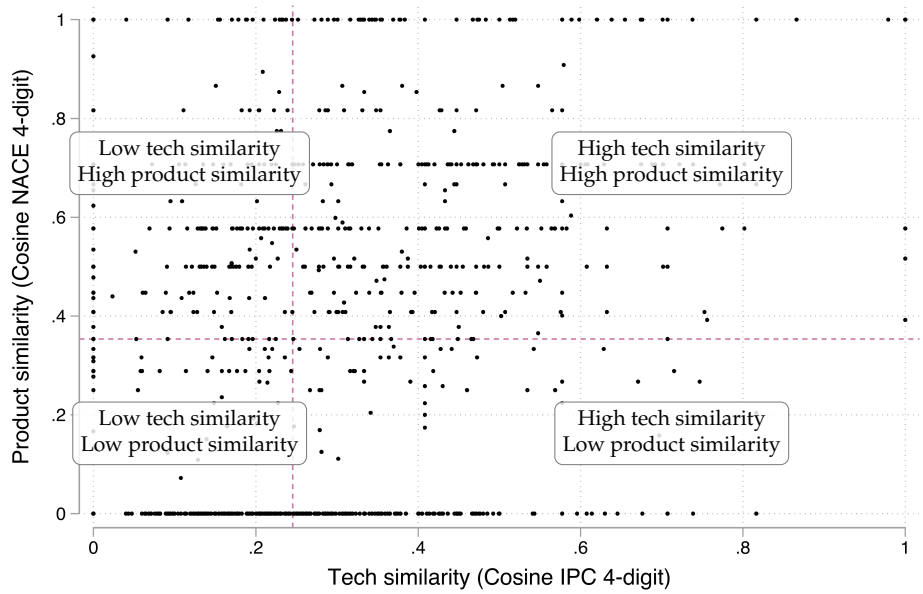
where $PatStock_{it}^{\text{target}}$ is the IHS-transformed stock of patents accumulated through acquisitions and depreciated using the perpetual inventory method. The annual acquisition indicator $M\&A^{NA}{}_{it}$ and fixed effects are defined as in equation (2). The coefficient of interest, β_2, captures the association between the acquired patent stock and markups conditional on the annual acquisition indicator and fixed effects.

Second, we exploit variation in the technological and product-market proximity between the merging parties to further investigate whether the evidence supports the innovation-shielding channel. Distinguishing these two dimensions matters because technological spillovers and product-market rivalry can have opposing effects (Bloom et al., 2013), while post-acquisition R&D outcomes depend on the complementarity of research activities and the rivalry between merging firms (Cassiman et al., 2005). To this end, we measure technological and market proximity between the acquirer and the target at the time of the deal as the cosine similarity of their four-digit IPC patent class and four-digit NACE sector vectors, respectively. The detailed construction of technological similarity is described in Section IB; product-market similarity is built analogously using NACE codes. For multiple acquisitions in a year, we average similarity across targets. We classify each acquirer using its maximum observed annual similarity in each dimension and split these measures at their respective medians in the panel of ever-treated firms. Figure 1 plots these measures.[18]

[17] We do not observe the individual patent rights transferred in each transaction and use the target's patent portfolio as a proxy.

[18] The corresponding plot and estimates using Jaccard similarity are reported in Supplemental Ap-

Figure 1: Technological and Market Proximity of Technological Acquirers (Cosine Similarity)



Notes: Each point represents an acquirer with observed similarity measures. The axes show the maximum observed annual cosine similarity in technological space (four-digit IPC classes) and product-market space (four-digit NACE sectors), calculated separately across its acquisitions. Same-year acquisitions are averaged across targets. The lines show the medians calculated over firm-year observations of ever-treated firms and define the four proximity groups.

Technological and product-market overlap between merging firms can favor the consolidation of patent rights around the acquirer's core technological domain and the removal of actual or potential rivals, strengthening its bargaining position and reducing competitive pressure (Hall and Ziedonis, 2001; Shapiro, 2000). Some technological relatedness can also facilitate the integration of acquired knowledge. Arts et al. (2025) discuss both integration and reduced rivalry in explaining why firms with unique technologies attract technologically proximate acquirers and close product-market rivals. Evidence of increased patenting following some early-stage pharmaceutical acquisitions close in both dimensions further illustrates the potential for synergies (Malek et al., 2025). These competing possibilities motivate interpreting proximity together with the evidence on patenting, productivity, and cost structure.

Following the baseline difference-in-differences design described in Section IIA, we compare the markup increase in the High-High (HH) group with that in the remaining groups by interacting the absorbing treatment indicator with indicators for these two groups; firms not conducting technological M&As form the reference category. We examine the four groups separately in Supplemental Appendix Table D6. Larger gains among acquirers with high proximity would be consistent with consolidation

pendix Figure D6 and Table D5.

strengthening their position in the technological domain.

III Results

A. Effect of Technological M&A on Firm Markup

Table 2 shows difference-in-differences estimates of the effect of technological M&As on firm markups across three alternative comparison groups. Column (1) uses the largest comparison group, comprising not-yet-treated firms and firms engaging in a non-technological M&A both before and after their own deal. Column (2) drops general-M&A firms observed before their deal, while Column (3) keeps only not-yet-treated firms. Because acquirers and non-acquirers differ markedly in size, R&D intensity, and profitability (Bena and Li, 2014; David, 2021; Nindl et al., 2024), we never compare M&A-active firms with firms that never undertake any M&A.

Across all three specifications, the coefficient on the technological-M&A indicator is positive, statistically significant, and similar across specifications (ranging from 0.0154 to 0.0184). Relative to the pre-treatment average markup of 1.386, this corresponds to an increase in markup of between 1.1 and 1.3 percent. Because the estimates are robust to progressively narrowing the comparison group, for the remainder of the analysis we adopt the specification with the largest control group (Column (1)).

Figure 2 reports the event-study estimates of model (1). The pre-treatment coefficients are close to zero and statistically insignificant. The joint test also does not reject the null that the pre-treatment coefficients are jointly zero ($p = 0.253$ in the TWFE specification), consistent with the parallel-trends assumption. Following the acquisition, markups rise and remain higher throughout the observed post-acquisition period: all six post-treatment coefficients are positive and significant at conventional levels, with somewhat larger magnitudes in later years. The widening confidence intervals at longer horizons reflect the smaller number of firms observed far from treatment.

To address the potential bias of the TWFE estimator under heterogeneous treatment timing (Goodman-Bacon, 2021), Figure 2 also reports results using the estimators of Sun and Abraham (2021) and Callaway and Sant'Anna (2021). The three estimators deliver similar average post-treatment effects: 0.018 (TWFE), 0.020 (Sun and Abraham, 2021), and 0.024 (Callaway and Sant'Anna, 2021), that is, between 1.3 and 1.7 percent of the pre-treatment mean markup.[19] We use the Callaway and Sant'Anna (2021) estimate as the reference in the remainder of the paper.

As indicated in Section IIA, the baseline results presented in this section use an absorbing treatment definition. Although switching is relatively uncommon in our

[19]The full set of event-study coefficients behind Figure 2, for all three estimators, is reported in Supplemental Appendix Table E1.

Table 2: Effect of Technological M&A on Markup

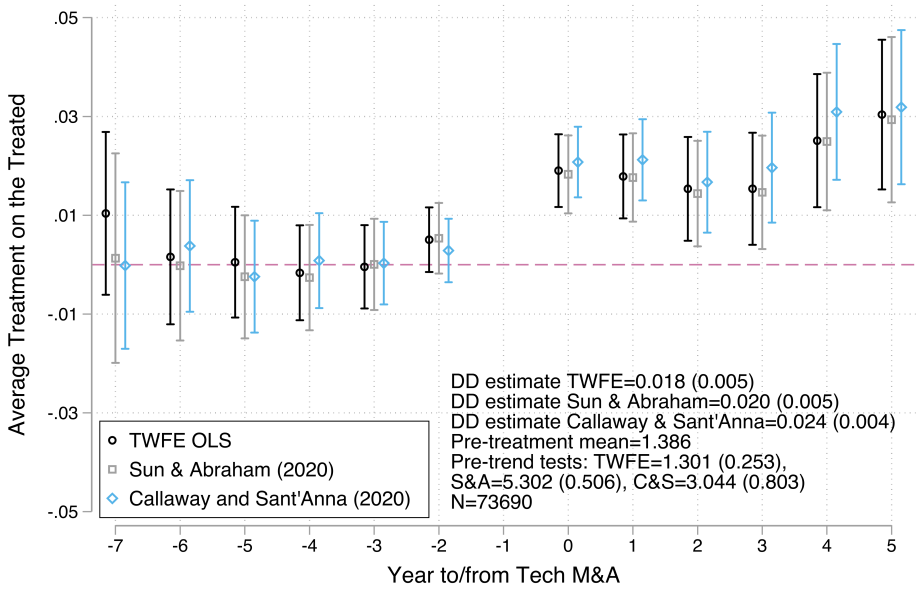
	(1)	(2)	(3)
M&A tech	0.0178***	0.0184***	0.0154***
	(0.0044)	(0.0046)	(0.0045)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
Control Group	NY + Pre + Post	NY + Post	NY
F-test	16.125	16.356	11.902
R sq	0.835	0.859	0.843
N (firm X year)	73690	48781	13356

Notes: The table reports TWFE difference-in-differences estimates of the effect of a technological M&A on firm markup, 2009–2020. The treatment indicator equals one from the year of the first technological M&A of the firm onward. The comparison group comprises not-yet-treated firms and firms with no technological M&A observed both before and after their own acquisitions (Column (1)); the same group excluding observations of firms with no technological M&A before their first acquisition (Column (2)); and not-yet-treated firms only (Column (3)). All columns include firm and industry-by-year fixed effects. Standard errors in parentheses are clustered at the firm level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

sample, treated firms may switch into and out of treatment during the observation period, which can pose problems for standard event-study estimators such as those in Figure 2. To assess robustness to this alternative treatment definition, Supplemental Appendix Figure D5 re-estimates the baseline markup equation using the estimator of De Chaisemartin and d'Haultfoeuille (2024), which aggregates specific 2×2 cells comparing "switchers" to "stayers" (i.e., units whose treatment remains unchanged) and is designed for settings with reversible treatments, staggered timing, and heterogeneous treatment effects. The estimates are in line with those obtained under the absorbing treatment definition in Figure 2.

The estimates presented here do not include covariates as controls (only firm and industry-by-year fixed effects) in light of recent econometric work on the use of covariates in difference-in-differences designs (Baker et al., 2026; Caetano and Callaway, 2024). Supplemental Appendix D3 discusses this issue and reports a set of alternative estimates obtained by including pre-treatment covariates to assess the robustness of the baseline specification to conditioning on relevant firm characteristics that differ between treated and comparison firms, as documented in Table 1 (Supplemental Appendix Figure D1). Supplemental Appendix D3 also includes re-estimations using an alternative comparison group constructed by matching on two-digit NACE indus-

Figure 2: Effect of Technological M&A on Markup



Notes: The figure plots event-study estimates of the effect of a technological M&A on firm markup from TWFE, Sun and Abraham (2021), and Callaway and Sant'Anna (2021), for up to seven years before and five years after the deal; the year before the deal is the reference period. The TWFE and Sun and Abraham (2021) specifications include firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A. Confidence intervals are at the 95 percent level and based on standard errors clustered at the firm level.

try[20] and acquisition year (Supplemental Appendix Figures D2 and D3). Finally, we re-estimate the event study using only not-yet-treated firms as the comparison group (Supplemental Appendix Figure D4), for which pre-treatment characteristics are more closely balanced (Supplemental Appendix Table C3). The baseline results are robust to all these alternative specifications, with average post-treatment ATTs between 0.021 and 0.023, against 0.024 in the baseline.

B. Heterogeneity in Markup Effects

1. Effect of Technological M&A by Firm Type, Industry, and Region

The overall effect illustrated above masks substantial heterogeneity across firms, industries, and regions. The aggregate rise in markups reflects a reallocation of activity toward high-markup firms and the growth of superstar firms (Autor et al., 2020; De Loecker et al., 2020). We investigate whether the impact of technological M&A on markups is driven by technologically dominant firms. We do so by distinguishing

[20]Matching on industry is particularly relevant because the concentration of technological acquirers in manufacturing accounts for part of the pre-treatment markup gap (Section IC).

between top and non-top R&D investors, defining top R&D investors as those firms in our sample that appear at least once among the world top 2,500 R&D investors in the EU Industrial R&D Investment Scoreboard over the period 2003–2024 (Nindl et al., 2024).[21]

Figure 3a shows the results. The average post-treatment ATT for top R&D investors is 0.036, a 2.6 percent increase relative to their pre-treatment mean markup (1.383), about 1.5 times the baseline effect of 1.7 percent (Section IIIA), while the effect for other firms is 0.012, or 0.8 percent of their pre-treatment mean (1.388). The difference between the two groups is statistically significant ($p = 0.010$). This suggests that the results are driven by the largest R&D investors, in line with the findings of the superstar literature (Autor et al., 2020; Tambe et al., 2020), which shows that already-dominant firms command higher market power. Technological M&As can thus be seen as one additional channel contributing to this consolidation.

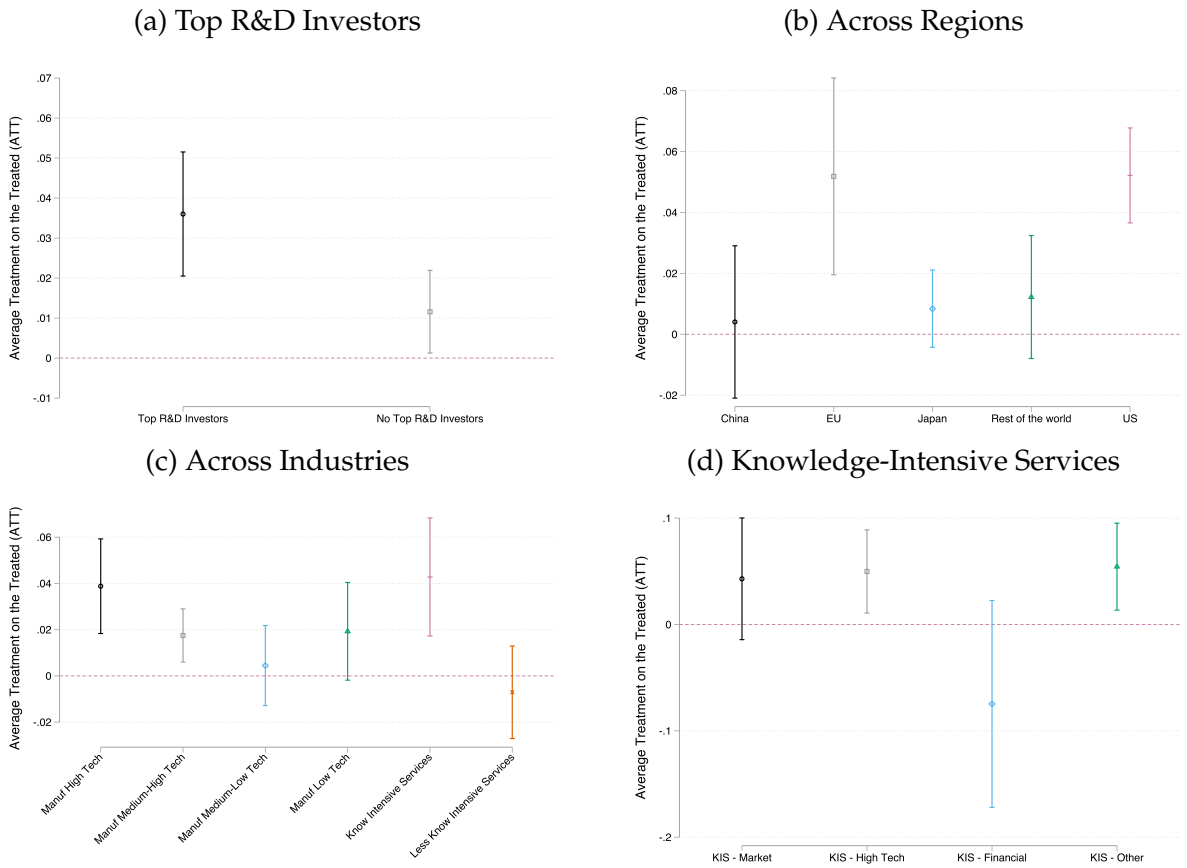
Figure 3b shows the estimated ATT for firms headquartered in five different world regions. The markup increase is concentrated in the United States and the European Union, where the effects are of similar magnitude: 0.052 in both, or 3.7 and 3.5 percent of the respective pre-treatment means (1.419 and 1.473). The point estimates for Japan (0.008), China (0.004), and the rest of the world (0.012) are small and statistically insignificant. This regional pattern is consistent with cross-country evidence on aggregate markup trends documenting stronger growth for the United States and Europe and showing more moderate increases or declines in Japan and China (De Loecker and Eeckhout, 2018).

Results across industries are shown in Figure 3c, adopting the Eurostat classification of industries by technological and knowledge intensity (Eurostat, 2014). The effect is positive and significant in high-tech manufacturing (0.039, or 2.7 percent), medium-high-tech manufacturing (0.018, or 1.3 percent), and knowledge-intensive services (the largest at 0.043, or 2.9 percent), while medium-low-tech manufacturing and less knowledge-intensive services display small and insignificant effects. Low-tech manufacturing displays a positive effect significant only at the 10 percent level (0.019). Figure 3d further disaggregates knowledge-intensive services, showing that the effect is driven mainly by high-tech and other knowledge-intensive services.[22] The stronger effects in high-tech manufacturing and knowledge-intensive services are consistent with the presence in these sectors of many superstar firms and major technology conglomerates, which derive competitive advantages from proprietary digital capital and

[21]The EU Industrial R&D Investment Scoreboard is an annual publication by the European Commission that, since 2004, monitors and benchmarks the R&D expenditures of the world's leading industrial investors.

[22]Other knowledge-intensive services comprise NACE Rev. 2 divisions 58 (publishing), 75 (veterinary activities), and 84–93 (public administration, education, health and social work, and arts and recreation) (Eurostat, 2014). Of the 65 technological acquirers in this group, 47 (72 percent) are in publishing, which includes software publishing, 10 in human health, and 4 each in education and in sports and recreation.

Figure 3: Effect of Technological M&A by Firm Type, Region, and Industry



Notes: Each panel reports the estimated effect of a technological M&A on firm markup, split along a different dimension of heterogeneity: (a) top R&D investors versus other firms; (b) world regions; (c) major industry aggregates; and (d) the four knowledge-intensive service categories. Top R&D investors are firms that appeared at least once in the list of the top 2,500 R&D investors over the period 2003–2024 (Nindl et al., 2024). Industry aggregation follows the Eurostat sectoral approach (Eurostat, 2014). All estimates are the average post-treatment ATTs from the Callaway and Sant'Anna (2021) estimator. Markups are pre-residualized on industry-by-year means of firms with no technological M&A. Confidence intervals are at the 95 percent level and based on standard errors clustered at the firm level.

intangible assets (Crouzet et al., 2022; Tambe et al., 2020) and may use acquired patent portfolios strategically to limit rivals' freedom to operate (Czarnitzki et al., 2020). This interpretation is consistent with evidence linking proprietary information technology to rising industry concentration (Bessen, 2020) and digital intensity to higher markups (Calligaris et al., 2024).

2. Technological Heterogeneity of Treatment

So far, we have analyzed the effect of technological M&As on markups, focusing on whether the acquisition of patent-filing firms increases the acquirer's market power. Patents differ, however, both in the value of the invention they protect and in the type of technology involved. We therefore conduct a treatment heterogeneity exercise to assess whether the results are robust to restricting the treatment to acquisitions involving IP5 patent families and whether substantial differences arise when technological M&As involve strategic technologies. To this end, we re-estimate model (1), redefining the treatment to include only acquisitions involving specific categories of patents held by the target at the time of the deal. As before, the comparison group comprises non-technological acquirers and eligible not-yet-treated firms, but firms undertaking a technological M&A of a different type from the one under scrutiny are excluded from the respective estimate.

First, we examine the robustness of our findings when restricting the treatment to acquisitions involving IP5 patent families, as a proxy for patent value (Dechezleprêtre et al., 2017; Harhoff et al., 2003).[23] Since each additional filing entails a cost, a family filed in at least two offices, one of them a major office, signals that the applicant expects the value of the invention to exceed the cost of multi-jurisdictional protection (Dechezleprêtre et al., 2017). In this case, the treatment is redefined as the acquisition of a target that holds at least one IP5 patent in stock at the time of the deal.

Results are shown in Figure 4. The average post-treatment ATT is 0.025, or about 1.8 percent relative to the pre-treatment markup of this group (1.408), broadly in line with the baseline result, which suggests that the effect does not depend on the acquisition of low-value patents.

Second, the effect of a technological M&A on markup may depend not only on the value of the patents involved but also on the type of technologies to which these patents relate. Beyond their intrinsic value, patents in specific technological domains may carry particular strategic importance. Concentrating the ownership of frontier technologies and technologies that play a central role within the current techno-economic paradigm can provide firms with a significant competitive advantage. For example, Petralia (2020) shows that firms with a higher number (or share) of general-purpose technology (GPT)

[23] Following Dernis et al. (2015), IP5 families are defined as those that include patents filed in at least two patent offices worldwide, including one of the five largest: CNIPA, EPO, JPO, KIPO, and USPTO.

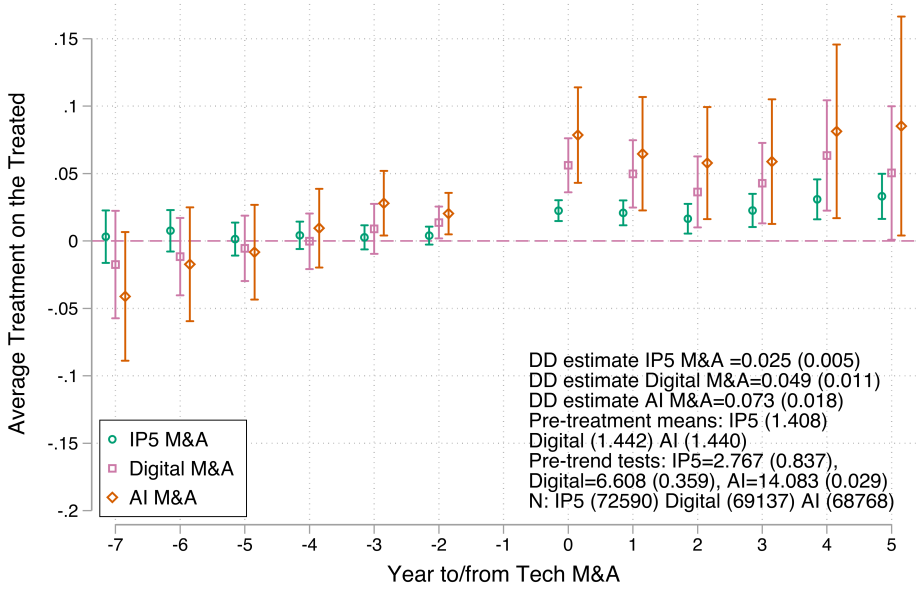
patents outperform others in terms of operating revenues. In the current technological paradigm, digital technologies—and increasingly artificial intelligence (AI)—are often discussed as exhibiting a general-purpose nature (Bresnahan and Trajtenberg, 1995). Their strategic relevance also raises potential policy concerns, as the concentration of frontier innovation in the hands of a few firms may have implications for both competition and the direction of technological change. Prior research links proprietary information and communication technologies (ICTs) to rising industry concentration (Bessen, 2020) and software investment and ICT patenting to higher firm markups (Calligaris et al., 2024).

In light of these considerations, we further refine our treatment definition by distinguishing technological M&As involving digital and AI technologies. We classify an acquisition as a digital (or AI) technological M&A when the acquired subsidiary holds at least one patent in the relevant technological domain at the time of the deal. The construction of the digital and AI indicators from the WIPO classification is described in Section IB and Supplemental Appendix B2; the digital indicator covers the dGPT categories other than AI.

Figure 4 shows the corresponding event-study estimates. The average post-treatment ATT for acquisitions involving digital patents is 0.049, or about 3.4 percent relative to this group's pre-treatment markup (1.442), substantially larger than the baseline. For AI-related acquisitions the effect is even larger, at 0.073 (about 5.1 percent relative to a pre-treatment markup of 1.440). Both average estimates are statistically significant. The results are consistent with evidence linking software investment and ICT patenting to higher firm markups (Calligaris et al., 2024) and in line with findings that the ownership of frontier technologies confers a substantial competitive advantage.

Given that the number of treated firms in the AI subgroup is small, the dynamic post-treatment estimates lose precision at longer horizons as the effective comparison group shrinks. Moreover, while the joint pre-treatment tests do not reject the null of no pre-treatment effects for IP5 and digital acquisitions, they reject it for AI-related acquisitions ($p = 0.029$). The AI estimate should therefore be interpreted as suggestive. Overall, the results suggest larger markup gains following acquisitions involving digital and AI patents, with potential implications not only for market power but also for the concentration of frontier innovation.

Figure 4: Effect of IP5, dGPT and AI Technological M&As on Markup



Notes: The figure shows the estimated effects of a technological M&A on firm markup distinguishing between acquisitions involving targets holding at least one IP5, digital, or AI patent, respectively. dGPT and AI patents are identified following the procedure in WIPO (2022) and described in Section IB and Supplemental Appendix B2. The figure plots the point estimates and confidence intervals from the estimator of Callaway and Sant'Anna (2021). Markups are pre-residualized on industry-by-year means of eligible firms never treated in the relevant category. Confidence intervals are at the 95 percent level and based on standard errors clustered at the firm level.

C. Mechanism Tests: Results

We next examine the mechanisms underlying the markup gains following technological M&As, using the tests outlined in Section IIB. We first assess the evidence for innovation enhancement and efficiency gains, before examining whether acquired patent portfolios and proximity between merging firms support an innovation-shielding interpretation.

1. Innovation-Enhancing Channel

Post-Acquisition Patenting Activity.—Table 3 reports the effect of technological M&As on annual patent flow. In the Callaway and Sant'Anna (2021) specification, technological acquirers file on average 8.2 more patents per year after the acquisition than comparison firms, while the Poisson estimate implies 5.7 percent higher expected annual patenting; neither estimate is statistically significant. Estimates on the markup estimation sample are similar (7.9 additional patents per year and 3.4 percent, respectively). The absence of a statistically detectable increase provides little support for increased inventive output as an explanation for the markup gains.

Table 4 reports the estimates of equation (2). Across the contemporaneous and subsequent patenting windows, the interaction coefficient $\hat{\beta}_2$ is close to zero and statistically insignificant. We therefore find no stronger association between markups in technological acquisition years and patenting in the deal year or the following one to three years.

Post-acquisition reorganization can offset the potential R&D synergies from combining knowledge bases. Ornaghi and Cassi (2026) document higher inventor exit and lower output among the inventors who stay at acquired pharmaceutical firms, while Szücs (2014) finds weaker R&D effort at targets and a decline in acquirers' R&D intensity associated with sales growth. These findings help explain why the acquisition of technological assets may not translate into an increase in inventive activity.

Productivity.—Figure 5 reports the productivity estimates from the design described in Section IIB. The pre-treatment coefficients are close to zero and statistically insignificant, consistent with the parallel-trends assumption for this outcome. Post-acquisition, the estimates indicate a modest decline in measured TFP of 1.7 percent with TWFE, 1.1 percent with Sun and Abraham (2021), and 1.9 percent with Callaway and Sant'Anna (2021). This pattern provides no support for operational efficiency gains as an explanation for the markup increase over the observed post-acquisition period. The initial decline may reflect operational disruption during the integration of acquired technologies and activities.

Fixed Costs.—Table 5 reports the fixed-cost-share estimates from the design described in Section IIB. Across specifications, we find limited evidence of an increase in the fixed cost share following a technological M&A. In Column (1), which excludes revenue

Table 3: Effect of Technological M&A on Patenting Activity

	(1) Patent flow	(2) Patent flow
M&A tech	8.2299	
	(7.9303)	
M&A tech		0.0559
		(0.1165)
Firm FE	Yes	Yes
Ind X Year FE	Yes	Yes
Model	C&S	Poisson HDFE
N (firm X year)	76781	33828

Notes: The table reports estimates of the effect of technological M&As on the merged entity's annual patent flow. Treatment is absorbing from the first technological acquisition. Column (1) reports the Callaway and Sant'Anna (2021) average post-treatment ATT, with not-yet-treated firms and firms with no technological M&A as comparison units, controlling for the pre-treatment patent stock of the acquirer; patent flow is pre-residualized on industry-by-year means of firms with no technological M&A. Column (2) reports the Poisson coefficient on post-acquisition status with firm and industry-by-year fixed effects; firms that never patent do not contribute to this estimate. Patent flow is untransformed in both columns. Estimates use all firm-years with patent data, including those without markup information. Standard errors in parentheses are clustered at the firm level. $^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.

Table 4: Patent Flow and Markup after Technological M&A

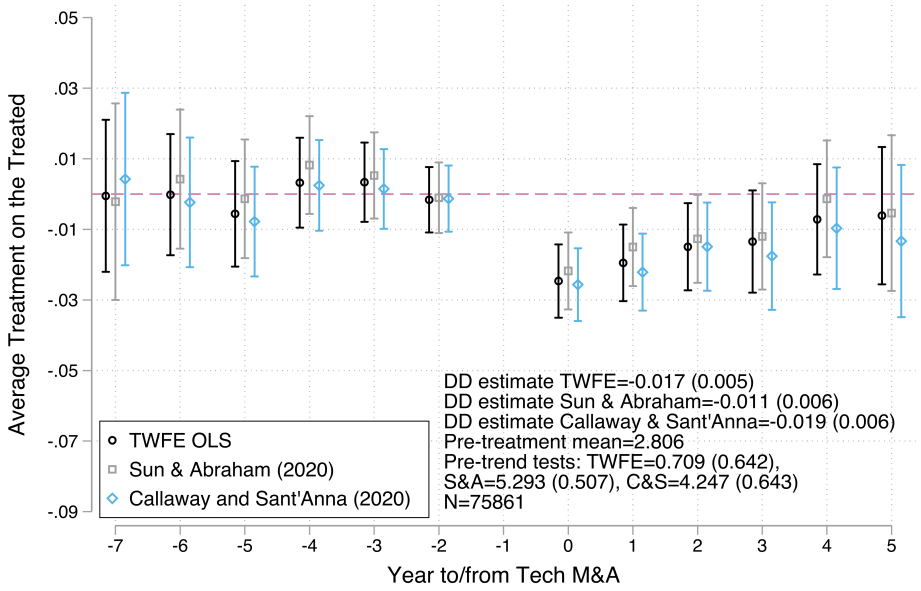
	(1)	(2)	(3)	(4)
M&A tech	0.0164***	0.0175***	0.0176***	0.0181***
	(0.0046)	(0.0047)	(0.0047)	(0.0048)
M&A tech X Pat flow	0.0010			
	(0.0017)			
M&A tech X Pat flow (t+1)		0.0004		
		(0.0018)		
M&A tech X Pat flow (t+2)			0.0004	
			(0.0018)	
M&A tech X Pat flow (t+3)				0.0001
				(0.0018)
Pat flow	-0.0007			
	(0.0013)			
Pat flow (t+1)		-0.0017		
		(0.0018)		
Pat flow (t+2)			-0.0021	
			(0.0022)	
Pat flow (t+3)				-0.0019
				(0.0024)
Firm FE	Yes	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes	Yes
F-test	9.905	10.179	10.292	9.989
R sq	0.835	0.835	0.835	0.835
N (firm X year)	73690	73690	73690	73690

Notes: The table reports estimates of equation (2). The dependent variable is firm markup. The treatment indicator equals one in each technological acquisition year (non-absorbing). Patent flow is measured in the deal year (Column (1)) and as its average over the deal year and the following one, two, or three years (Columns (2)–(4)), and is inverse-hyperbolic-sine transformed. All columns include firm and industry-by-year fixed effects. Standard errors in parentheses are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

controls, the estimate is close to zero and statistically insignificant (0.1 percentage points of revenue). Adding revenues as a control in Columns (2) and (3) yields larger positive estimates for both TWFE (0.8 percentage points) and Callaway and Sant'Anna (2021) (0.7 percentage points), but neither is statistically significant at the 5 percent level. While De Ridder (2024) shows that intangible-driven changes in cost structure can explain a substantial part of the aggregate rise in markups, these results do not support the hypothesis that markup gains following technological M&As are driven by a shift toward scalable, fixed-cost-intensive production structures.

Additional estimates of the effect of technological M&As on profitability and excess markup are reported in Supplemental Appendix D2. These estimates also indicate higher operating profits following technological M&As (3.0 percent in the Callaway and

Figure 5: Effect of Technological M&A on Productivity (Event Study)



Notes: The figure plots event-study estimates of the effect of a technological M&A on log TFP (Ackerberg et al., 2015) from TWFE, Sun and Abraham (2021), and Callaway and Sant'Anna (2021), for up to seven years before and five years after the deal, relative to the year before the deal. The TWFE and Sun and Abraham (2021) specifications include firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, TFP is pre-residualized on industry-by-year means of firms with no technological M&A. Bars show 95 percent confidence intervals based on standard errors clustered at the firm level.

Sant'Anna (2021) estimate), supporting a market-power interpretation of the markup increase. The positive and statistically significant effect on excess markup (0.010, or 1.0 percentage point of revenue; Supplemental Appendix Table D4) shows that the markup increase persists after adjustment for the industry-by-year fixed-cost benchmark, suggesting that the gains reflect more than the need to cover higher fixed costs at the industry level. These findings are consistent with De Loecker et al. (2020), who show that rising overhead costs do not fully account for the increase in US markups between 1980 and 2016.

Taken together, the absence of a detectable increase in patenting, the decline in measured productivity, and the limited evidence of a higher fixed cost share provide little support for R&D synergies or efficiency gains as explanations for the markup increase. We next examine whether the evidence is more consistent with the strategic consolidation of existing technological assets.

Table 5: Effect of Technological M&A on Fixed Cost Share

	(1)	(2)	(3)
M&A tech	0.0012	0.0076*	
	(0.0039)	(0.0040)	
M&A tech			0.0069
			(0.0043)
Estimator	TWFE	TWFE	CS
Revenues	No	ihs	ihs
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
F-test	0.088	17.077	
R sq	0.543	0.546	
N (firm X year)	68277	68277	67997

Notes: The table reports estimates of the effect of technological M&A on the acquirer's fixed cost share, defined as the difference between the Lerner index and the operating margin. Column (1) includes no controls; Columns (2) and (3) control for IHS-transformed revenues to account for size-related differences in fixed cost levels. Columns (1) and (2) report TWFE estimates; Column (3) reports Callaway and Sant'Anna (2021) estimates. All specifications include firm and industry-by-year fixed effects. Standard errors in parentheses are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2. Innovation-Shielding Channel

Acquired Patent Stock.—We next examine the association between markups and accumulated acquired patent portfolios using equation (3), described in Section IIB2.

Table 6 reports the estimates. In Column (1), the acquired patent stock is positively and significantly associated with acquirer markups ($\hat{\beta}_2 = 0.0072$): doubling the acquired stock is associated with a markup increase of about 0.005 ($0.0072 \times \ln 2$), or 0.4 percent of the pre-treatment mean markup (1.386). While the binary treatment identifies acquisitions of targets with at least one patent, this result highlights the intensive margin of patent accumulation: the acquisition of a larger patent portfolio is associated with higher markups, conditional on firm and industry-by-year fixed effects and the annual acquisition indicator. The coefficient on the latter is small and statistically insignificant ($\hat{\beta}_1 = 0.0059$).

Column (2) adds patent flow alongside patent stock. The stock coefficient is virtually unchanged ($\hat{\beta}_2 = 0.0073$), while the flow coefficient is small and statistically indistinguishable from zero ($\hat{\beta} = -0.0011$). Together with the patent-outcome estimates

Table 6: Patent Stock and Markup after Technological M&A

	(1) Stock	(2) Flow + Stock
M&A tech	0.0059	0.0059
	(0.0037)	(0.0037)
Pat flow		-0.0011
		(0.0014)
Acquired pat. stock (PIM)	0.0072***	0.0073***
	(0.0020)	(0.0020)
Firm FE	Yes	Yes
Ind X Year FE	Yes	Yes
F-test	15.035	10.468
R sq	0.846	0.846
N (firm X year)	67492	67492

Notes: The table reports the association between the acquired patent stock and firm markup. The dependent variable is firm markup. Column (1) includes the acquired patent stock alongside the M&A indicator; Column (2) additionally controls for patent flow. Patent stock is constructed using the perpetual inventory method (PIM) and IHS-transformed. The treatment indicator equals one in each technological acquisition year. All specifications include firm and industry-by-year fixed effects. Standard errors in parentheses are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

above, these results suggest that the markup gains are linked more closely to the accumulation of existing patent portfolios than to increased inventive activity. This pattern is consistent with an innovation-shielding interpretation, in which the consolidation of patent rights strengthens the acquirer's bargaining position vis-à-vis rivals.

This interpretation connects with Akcigit and Ates (2021, 2023), whose frameworks link weaker knowledge diffusion to higher markups and whose empirical evidence documents increasing concentration of patent ownership. Our results identify technological acquisitions as a route through which firms accumulate patent portfolios associated with higher markups.[24]

Technological and Market Proximity.—We compare post-acquisition markup increases between the High-High (HH) group and the remaining groups (HL: high technological and low product-market proximity; LH: low technological and high product-market proximity; LL: low proximity in both dimensions), using the firm-level classification described in Section IIB2.

[24]This shielding interpretation may appear to be at odds with evidence that technological acquisitions can raise the forward citations received by the target's patents (Yang, 2025). The two pieces of evidence may instead be reconciled by noting that, in a strategic-patenting environment, additional citations may arise as rivals seek to invent around acquired patents. Increased citations can coexist with consolidated patent ownership strengthening the acquirer's bargaining position vis-à-vis rivals.

Table 7 shows that markup gains following technological M&As are larger in the HH group, characterized by high technological and product-market proximity. The Callaway and Sant'Anna (2021) estimates in Columns (2) and (3) yield ATTs of 0.050 for the HH group and 0.017 for the remaining groups, that is, 3.5 and 1.3 percent of their respective pre-treatment means (1.414 and 1.355); the difference is statistically significant ($p = 0.001$). The TWFE estimates in Column (1) confirm this pattern (0.039 and 0.010; $p = 0.006$).

These larger gains are consistent with consolidation strengthening the acquirer's market position where technological and product-market overlap is substantial. Combining overlapping patent portfolios can strengthen bargaining positions, while acquisitions involving firms operating in similar sectors can reduce competitive pressure. Together with the evidence on acquired patent stock and post-acquisition patenting, this pattern is consistent with an innovation-shielding interpretation of the markup increase.

Supplemental Appendix Table D5 replicates the analysis using Jaccard similarity instead of cosine similarity and yields similar estimates, so the pattern is robust to this alternative measure.[25] The full four-quadrant decomposition in Supplemental Appendix D6 provides further detail (Table D6): larger markup gains occur in both groups with high technological proximity, including when product-market proximity is low (HL). Although the HH group has the largest point estimate, its effect is not statistically distinguishable from that of HL. The low-technological-proximity groups display smaller and less consistently significant effects.

Taken together, these results suggest that technological overlap is central to the markup gains following technological M&As. The sizable gains in both HH and HL are consistent with the consolidation of related patent portfolios strengthening the acquirer's position, including where product-market overlap is limited. This concentration along the technological dimension connects to Haucap et al. (2019), who find that, following horizontal pharmaceutical mergers, declines in patenting are concentrated in technology classes where the merging firms' and rivals' innovation activities overlapped before the deal. Read alongside the positive association between acquired patent stocks and markups and the limited evidence of increased inventive output or efficiency gains, these findings are consistent with an innovation-shielding interpretation. Whether consolidation strengthens the acquirer's position through a denser patent portfolio, the removal of an actual or potential rival, or a combination of the two remains an open question.

[25] Supplemental Appendix B3 defines both similarity measures and explains their different normalizations.

Table 7: Technological M&A and Proximity: HH vs Rest (Cosine Similarity)

	(1) TWFE	(2) CS: HH	(3) CS: Rest
M&A tech × HH	0.0388***		
	(0.0091)		
M&A tech × Rest (HL+LH+LL)	0.0101*		
	(0.0054)		
M&A tech		0.0502***	0.0172***
		(0.0081)	(0.0052)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
F: HH=Rest	7.571		
p: HH=Rest (TWFE)	0.006		
z: HH=Rest		3.416	
p: HH=Rest (CS z-test)		0.001	
F-test (model)	10.473		
R sq	0.836		
N (firm X year)	73358	63420	67458

Notes: The table reports TWFE and Callaway and Sant'Anna (2021) estimates of the markup effect of technological M&As for acquirers in the HH group (high technological and high product-market proximity) and in the remaining groups (HL, LH, and LL). Proximity is the cosine similarity of the four-digit IPC and four-digit NACE vectors of acquirer and target, split at the median among ever-treated firms. The treatment indicator is absorbing. The table reports F-tests (TWFE) and z-tests (Callaway and Sant'Anna 2021, CS) of equal effects for the two groups. The TWFE specification includes firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors in parentheses are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

IV Quantification and Economic Relevance

The estimated coefficients can be given economic meaning in different ways. First, we take as reference the Callaway and Sant'Anna (2021) baseline estimate of 0.024 reported in Figure 2. Evaluated at the pre-treatment mean, the estimated effect raises the price-to-marginal-cost ratio from 1.386 to 1.410, a 1.7 percent increase. Equivalently, the price premium above marginal cost rises from 38.6 to 41.0 percent, that is an increase of 2.4 percentage points.

Second, in terms of the implied Lerner index ($L = 1 - 1/\mu$), the pre-treatment markup corresponds to about 27.8 cents of each unit of revenue accruing as margin over marginal cost. Conducting a technological M&A raises this to roughly 29.1 cents, an increase of about 1.2 percentage points of revenue.

The central question is whether these effects are economically relevant. A change of

0.024 is small relative to the cross-sectional dispersion of markups: in our estimation sample, the interquartile range of the markup is 0.35, from 1.23 to 1.58. Nevertheless, the effect is sizable as a causal within-firm effect. A natural benchmark is the secular trend: De Loecker et al. (2020) document an increase in the aggregate US markup from 1.21 in 1980 to 1.61 in 2016, an average of about 0.0111 per year. Our estimate is therefore equivalent to roughly 2.1 years of aggregate markup growth. In the same vein, Stiebale and Szücs (2022) find that the non-merging rivals of merging firms raise their markups by 2–4 percent on average after a merger; since our acquirer effect is additional to any such rival response, the market-wide markup impact of a technological merger combines the two. Together, these benchmarks indicate an economically relevant effect of technological M&As on markups.

Table 8 reports the quantified effects across the relevant model specifications in the paper. The effect scales systematically with the strategic value of the acquired technology and the market position of the acquirer. Compared to the baseline estimate (1.7 percent markup increase), it rises for firms, regions, and technologies where the innovation-shielding channel is expected to be strongest. It is about 1.5 times as large for top R&D investors (2.6 percent) and larger in the most technology- and knowledge-intensive sectors (2.7 percent in high-tech manufacturing, 2.9 percent in knowledge-intensive services) and among US- and EU-based acquirers (3.7 and 3.5 percent, respectively), which are the advanced economies that host the bulk of superstar firms. The effect is largest for acquisitions of frontier digital and AI technologies, where the markup rises by 3.4 and 5.1 percent, respectively, although the AI estimate should be taken as suggestive given that its pre-treatment trend test is rejected. In terms of the Lerner index, these effects imply an increase in the margin over marginal cost ranging from about 1.2 to 3.3 percentage points of revenue. Benchmarked against the secular trend of the US markup, the effects range from about two years of aggregate markup growth for the baseline to nearly seven years for the AI subsample.

Table 8: Quantification of the Markup Effect across Specifications

	ATT (CS)	Markup incr. (%)	Lerner (pp rev.)	Trend (years)
Baseline	0.024	1.7	1.2	2.1
By firm type				
Top R&D investors	0.036	2.6	1.8	3.2
By region				
United States	0.052	3.7	2.5	4.7
European Union	0.052	3.5	2.3	4.7
By sector				
High-tech manufacturing	0.039	2.7	1.9	3.5
Knowledge-intensive services	0.043	2.9	1.9	3.9
By acquired technology				
IP5 patents	0.025	1.8	1.2	2.3
Digital technologies	0.049	3.4	2.3	4.4
Artificial intelligence	0.073	5.1	3.3	6.6

Notes: The table reports the economic magnitude of the markup effect of technological M&As across specifications. ATT is the Callaway and Sant'Anna (2021) average post-treatment effect; all reported estimates are statistically significant at conventional levels. *Markup increase* is the ATT as a percentage of the pre-treatment mean markup of the treated firms in each group: 1.386 (baseline), 1.383 (top R&D investors), 1.419 (United States), 1.473 (European Union), 1.420 (high-tech manufacturing), 1.463 (knowledge-intensive services), 1.408 (IP5), 1.442 (digital), and 1.440 (AI). *Lerner increase* is the implied change in the Lerner index $L = 1 - 1/\mu$, i.e., the percentage-point increase in the share of revenue accruing as markup over marginal cost. *Aggregate-trend equivalent* is the ATT divided by the average annual increase in the aggregate US markup (0.0111 per year) documented by De Loecker et al. (2020). The AI estimate fails the pre-trend test ($p = 0.029$) and should be read as suggestive.

V Conclusions

We examine whether technological M&As, defined as acquisitions of targets holding at least one patent family, raise acquirer markups, and through which channels. Drawing on a global panel of 7,796 listed firms over 2009–2020 that links acquisition records and financial, ownership, and patent data from Orbis, we estimate the effect on acquirer markups with a difference-in-differences design that compares technological acquirers with acquirers that have not yet completed, or never complete, a technological M&A. We then distinguish an innovation-enhancing channel, through which acquired knowledge raises inventive output and efficiency, from an innovation-shielding channel, through which the acquired patent portfolio protects the acquirer from rivals.

Across industries and countries, technological M&As raise acquirer markups by 1.7 percent on average (1.3–1.7 percent across estimators), and the effect persists over the observed post-acquisition period. Patenting by the acquirer and its subsidiaries shows no statistically detectable increase, and measured productivity declines in the years immediately after the deal, so that markups rise without a detectable gain in innovation or efficiency. The acquisition of patent-filing firms is thus a specific corporate strategy through which firms raise their markups, and its effect is larger among top R&D investors (2.6 percent), among acquirers headquartered in the United States and the European Union, and in high-tech manufacturing and knowledge-intensive services. The point estimates are also larger for acquisitions involving digital and AI patents, although the AI estimate rests on few treated firms.

We find that markup gains increase with the size of the acquired patent stock and with the technological proximity between acquirer and target, and the largest point estimate arises when the target is also close in product-market space. Together with the absence of a shift toward fixed costs, these patterns are more consistent with the innovation-shielding channel than with innovation or efficiency gains, although the mechanism tests do not identify causal mediation. They are also consistent with evidence that strategic patents are associated with higher market concentration and with lower productivity and patenting of competitors (Veihl, 2022). Combined with the rising markups and, if anything, falling innovation among the rivals of merging firms documented by Stiebale and Szücs (2022), our results indicate that technological mergers can raise markups on both sides of the transaction without increasing innovation on the acquirer side. This pattern is consistent with the decline in knowledge diffusion modeled by Akcigit and Ates (2023). We find no efficiency or innovation gains that would offset the higher price-cost margin; absent such gains, a merger harms consumers in the model of Bourreau et al. (2026). These findings support merger review criteria that take innovation outcomes into account, not only static market shares.

Our estimates also inform the design of merger review in three respects.

To better assess the innovation effects of mergers, merger review processes could be recalibrated to integrate forward-looking innovation metrics—such as patent quality, citation indices, and R&D intensity—alongside traditional static assessments of market concentration. An "innovation defense" could allow firms to justify consolidation only when they can demonstrate verifiable R&D synergies *ex ante*, while *ex post* monitoring of cleared M&A deals would ensure that anticipated dynamic efficiencies materialize rather than simply strengthen market power. Such suggestions are in line with the Draghi Report's advocacy (Draghi, 2024) for a merger control framework that systematically incorporates innovation potential and permits rigorous post-merger evaluation, which would deter opportunistic "killer acquisitions" of nascent innovators (Cunningham et al., 2021). The case for *ex post* monitoring is reinforced by Stiebale and Szücs (2022), who find that remedies imposed by the European Commission reduced, but did not offset, post-merger markup increases.

Three limitations qualify these results and suggest directions for further research. First, technological M&As are defined by the patent holdings of targets, so acquisitions of innovative firms and start-ups that do not hold patents are not classified as technological, and the markup effects of acquiring unpatented technology are not measured. Second, while markups show no differential trends before the deal and the estimates are robust to accounting for observed differences between acquirers, selection of targets on an unobserved technological trajectory cannot be fully excluded; a comparison of completed with failed or withdrawn technological acquisitions, following Seru (2014), could provide a sharper counterfactual. Finally, the mechanism tests do not identify causal mediation and cannot separate the effect of a denser patent portfolio from that of removing an actual or potential rival; outcomes for rivals and at the market level would help distinguish the two.

References

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Akcigit, U. and Ates, S. T. (2021). Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics*, 13(1):257–298.

Akcigit, U. and Ates, S. T. (2023). What happened to us business dynamism? *Journal of Political Economy*, 131(8):2059–2124.

Andrews, D., Criscuolo, C., and Gal, P. (2016). The best versus the rest: The global productivity slowdown, divergence across firms and the role of public policy. OECD Productivity Working Papers 5, OECD, Paris.

Arts, S., Cassiman, B., and Hou, J. (2025). Technology differentiation, product market rivalry, and m&a transactions. *Strategic Management Journal*, 46(4):837–862.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly journal of economics*, 135(2):645–709.

Baker, A., Callaway, B., Cunningham, S., Goodman-Bacon, A., and Sant'Anna, P. H. (2026). Difference-in-differences designs: A practitioner's guide. *Journal of Economic Literature*, 64(2):498–557.

Bena, J. and Li, K. (2014). Corporate innovations and mergers and acquisitions. *The Journal of Finance*, 69(5):1923–1960.

Bessen, J. (2020). Information technology and industry concentration. *Journal of Law & Economics*, 63(3):531.

Bhattacharya, V., Illanes, G., and Stillerman, D. (2026). Merger effects and antitrust enforcement: Evidence from US consumer packaged goods. *American Economic Review*, 116(8):2996–3035.

Bloom, N., Schankerman, M., and Van Reenen, J. (2013). Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393.

Bourreau, M., Jullien, B., and Lefouili, Y. (2026). Horizontal mergers and incremental innovation. *The RAND Journal of Economics*. Early View, online 25 February 2026.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of econometrics*, 65(1):83–108.

Caetano, C. and Callaway, B. (2024). Difference-in-differences when parallel trends holds conditional on covariates. Technical report, arXiv preprint arXiv:2406.15288.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

Calligaris, S., Criscuolo, C., and Marcolin, L. (2024). Mark-ups in the digital era. Working paper, London School of Economics and Political Science, Centre for Economic Performance.

Cao, L. and Zhu, L. (2025). Mergers and acquisitions and the aggregate markup. *Macroeconomic Dynamics*, 29:e115.

Cassiman, B., Colombo, M. G., Garrone, P., and Veugelers, R. (2005). The impact of m&a on the r&d process: An empirical analysis of the role of technological-and market-relatedness. *Research policy*, 34(2):195–220.

Cohen, W. M., Levinthal, D. A., et al. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative science quarterly*, 35(1):128–152.

Crouzet, N., Eberly, J. C., Eisfeldt, A. L., and Papanikolaou, D. (2022). The economics of intangible capital. *Journal of Economic Perspectives*, 36(3):29–52.

Cunningham, C., Ederer, F., and Ma, S. (2021). Killer acquisitions. *Journal of political economy*, 129(3):649–702.

Czarnitzki, D., Hussinger, K., and Leten, B. (2020). How valuable are patent blocking strategies? *Review of Industrial Organization*, 56(3):409–434.

David, J. M. (2021). The aggregate implications of mergers and acquisitions. *The Review of Economic Studies*, 88(4):1796–1830.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *Econometrica*, 88(6):2191–2224.

De Chaisemartin, C. and d'Haultfoeuille, X. (2024). Difference-in-differences estimators of intertemporal treatment effects. *Review of Economics and Statistics*, pages 1–45.

De Loecker, J. and Eeckhout, J. (2018). Global market power. Working Paper 24768, National Bureau of Economic Research.

De Loecker, J. and Eeckhout, J. (2026). The macroeconomics of market power. *Annual Review of Economics*, 18:[to fill].

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644.

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

De Ridder, M. (2024). Market power and innovation in the intangible economy. *American Economic Review*, 114(1):199–251.

De Ridder, M., Grassi, B., and Morzenti, G. (2026). The hitchhiker's guide to markup estimation: assessing estimates from financial data. *Econometrica*, 94(1):137–168.

Dechezleprêtre, A., Ménière, Y., and Mohnen, M. (2017). International patent families: from application strategies to statistical indicators. *Scientometrics*, 111(2):793–828.

Dernis, H., Dosso, M., Hervás Soriano, F., Millot, V., Squicciarini, M., and Vezzani, A. (2015). World corporate top R&D investors: Innovation and IP bundles. EUR 27129 EN, A JRC and OECD common report, Publications Office of the European Union, Luxembourg.

Draghi, M. (2024). *The future of European Competitiveness: In-depth analysis and recommendations*. European Commission.

Ederer, F., Seibel, R., and Simcoe, T. (2026). Digital (killer?) acquisitions. Working Paper 35762, National Bureau of Economic Research.

Eurostat (2014). High-tech industry and knowledge-intensive services (htec): Reference metadata in euro sdmx metadata structure (esms). `https://ec.europa.eu/eurostat/cache/metadata/tmp/htec_esms_4d.htm#annex1391770094474`. Metadata last certified: 17 January 2014; accessed 16 June 2025.

Florance, K. (2012). Announcing the Netflix Open Connect Network. Netflix, June 4. `https://about.netflix.com/en/news/announcing-the-netflix-open-connect-network`. Accessed October 7, 2026.

Frakes, M. D. and Wasserman, M. F. (2025). Strategic patenting: Evidence from the biopharmaceutical industry. NBER Working Paper 34024, National Bureau of Economic Research.

Gal, P. (2013). Measuring total factor productivity at the firm level using oecd-orbis. OECD Economics Department Working Papers 1049, OECD, Paris.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of econometrics*, 225(2):254–277.

Hall, B. H., Jaffe, A. B., and Trajtenberg, M. (2005). Market value and patent citations. *The RAND Journal of Economics*, 36(1):16–38.

Hall, B. H. and Ziedonis, R. H. (2001). The patent paradox revisited: an empirical study of patenting in the us semiconductor industry, 1979-1995. *RAND Journal of Economics*, pages 101–128.

Harhoff, D., Scherer, F. M., and Vopel, K. (2003). Citations, family size, opposition and the value of patent rights. *Research Policy*, 32(8):1343–1363.

Haucap, J., Rasch, A., and Stiebale, J. (2019). How mergers affect innovation: Theory and evidence. *International Journal of Industrial Organization*, 63:283–325.

Iacus, S. M., King, G., and Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. *Political analysis*, 20(1):1–24.

Imbens, G. W. and Wooldridge, J. M. (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1):5–86.

Jin, G. Z., Leccese, M., and Wagman, L. (2023). How do top acquirers compare in technology mergers? new evidence from an s&p taxonomy. *International Journal of Industrial Organization*, 89:102891.

Kahn-Lang, A. and Lang, K. (2020). The promise and pitfalls of differences-in-differences: Reflections on 16 and pregnant and other applications. *Journal of Business & Economic Statistics*, 38(3):613–620.

Kalemli-Özcan, Ş., Sørensen, B. E., Villegas-Sanchez, C., Volosovych, V., and Yeşiltaş, S. (2019). How to construct nationally representative firm level data from the orbis global database: New facts and aggregate implications. Technical report, National Bureau of Economic Research.

Kamepalli, S. K., Rajan, R., and Zingales, L. (2020). Kill zone. Technical report, National Bureau of Economic Research.

KAYAK Software Corporation (2012). KAYAK reports record results and agrees to acquisition by Priceline.com Incorporated. Company press release, GlobeNewswire, November 8. https://www.globenewswire.com/news-release/2012/11/08/503772/10011755/en/KAYAK-Reports-Record-Results-and-Agrees-to-Acquisition-by-Priceline-com-Incorp.html. Accessed October 7, 2026.

Levonyan, V. and Mengano, P. (2024). Mergers and market power. Working Paper.

Malek, J., Seldeslachts, J., and Veugelers, R. (2025). Are m&as spurring or stifling innovation? evidence from antidiabetic drug development. Technical report, DIW Discussion Papers.

Miller, N. H. and Weinberg, M. C. (2017). Understanding the price effects of the MillerCoors joint venture. *Econometrica*, 85(6):1763–1791.

Netflix (2017). Netflix acquires Millarworld. Company press release, August 7. https://about.netflix.com/en/news/netflix-acquires-millarworld-1. Accessed October 7, 2026.

Nindl, E., Napolitano, L., Confraria, H., Rentocchini, F., Fako, P., Gavigan, J., and Tuebke, A. (2024). The 2024 eu industrial r&d investment scoreboard. Technical Report KJ-01-24-165-EN-N (online),KJ-01-24-165-EN-C (print), European Commission, Luxembourg (Luxembourg).

Ornaghi, C. and Cassi, L. (2026). Acquisitions, inventors' turnover, and innovation: Evidence from the pharmaceutical industry. *European Economic Review*. Article 105263.

Petralia, S. (2020). Mapping general purpose technologies with patent data. *Research Policy*, 49(7):104013.

Phillips, G. M. and Zhdanov, A. (2013). R&D and the incentives from merger and acquisition activity. *Review of Financial Studies*, 26(1):34–78.

Sant'Anna, P. H. and Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of econometrics*, 219(1):101–122.

Seru, A. (2014). Firm boundaries matter: Evidence from conglomerates and R&D activity. *Journal of Financial Economics*, 111(2):381–405.

Shapiro, C. (2000). Navigating the patent thicket: Cross licenses, patent pools, and standard setting. *Innovation policy and the economy*, 1:119–150.

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

Stiebale, J. and Szücs, F. (2026). Study on the impact of mergers reviewed by the European Commission on innovation and markups: An empirical investigation. Technical report, European Commission, Directorate-General for Competition. Publications Office of the European Union.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199.

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

Szücs, F. (2014). M&a and r&d: Asymmetric effects on acquirers and targets? *Research Policy*, 43(7):1264–1273.

Tambe, P., Hitt, L., Rock, D., and Brynjolfsson, E. (2020). Digital capital and superstar firms. Working Paper 28285, National Bureau of Economic Research.

Veihl, M. (2022). Strategic patenting: The dark side of patents. Working paper. SSRN Working Paper No. 3900690.

WIPO (2022). *World intellectual property report 2022: The direction of innovation*. WORLD INTELLECTUAL PROPERTY ORGANIZATION.

Wollmann, T. G. (2019). Stealth consolidation: Evidence from an amendment to the Hart-Scott-Rodino act. *American Economic Review: Insights*, 1(1):77–94.

Wooldridge, J. M. (2009). On estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters*, 104(3):112–114.

Yang, Z. (2025). Technology M&As and knowledge diffusion. Working paper, University of Southern California.

Supplemental Appendix A: Markups

A1 Theory of Markups

The empirical approach used to recover firm-level markups relies on standard cost minimization conditions for variable inputs free of adjustment costs and on the estimation of output elasticity. In particular, we assume firm i's production technology at time t is: $Q_{it} = Q_{it}(X_{it}^1, ..., X_{it}^V, K_{it}, \omega_{it})$ where V are variable inputs such as labor or intermediate inputs and ω is the productivity level. We assume the production function to be continuous and twice differentiable in its arguments and that firms employ labor, intermediate inputs, and capital to produce. Hence, they face the following cost minimization problem:

$$\min_{\{K_{it}, L_{it}, M_{it}\}} w_{it}L_{it} + r_{it}K_{it} + p_{it}^M M_{it}, \quad s.t. : Q_{it}(L_{it}, M_{it}, K_{it}) \leq Q_{it}, \tag{A1}$$

where w_{it} , r_{it} , and p_{it}^M denote the wage rate, the interest rate on capital, and the intermediate input price, respectively. We can derive the associated Lagrangian as:

$$\mathcal{L}(L_{it}, K_{it}, M_{it}, \lambda_{it}) = w_{it}L_{it} + r_{it}K_{it} + p_{it}^M M_{it} + \lambda(Q_{it} - Q_{it}(\cdot)) \tag{A2}$$

We then derive the first-order conditions for any variable inputs associated with the Lagrangian function such that:

$$\frac{\partial \mathcal{L}}{\partial X_{it}^v} = P_{it}^{X^v} - \lambda_{it}\frac{\partial Q_{it}(\cdot)}{\partial X_{it}^v} = 0 \tag{A3}$$

where the shadow price λ_{it} is the marginal cost of production at a given level of output. Rearranging terms, multiplying both sides by $\frac{X_{it}}{Q_{it}}$ and defining $\mu_{it} \equiv \frac{P_{it}}{\lambda_{it}}$ the following expression for markups can be derived:

$$\mu_{it} = \theta_{it}^X(\alpha_{it}^X)^{-1} \tag{A4}$$

where θ_{it}^X is the output elasticity of an input X and α_{it}^X is the share of expenditures on input X_{it} in total sales $(P_{it}Q_{it})$.

A2 Implementation of the De Loecker and Warzynski (2012) Method

Data Cleaning Before Markup Estimation

The De Loecker and Warzynski (2012) method leverages firm-level balance sheet data to infer markups indirectly. The balance sheet and income statement data recorded in Orbis is obtained from company accounts, therefore its coverage may differ across countries or across years. Before computing the markup estimates, we clean the data and impose homogenizing criteria. The cleaning steps we perform are closely in line with previous OECD analyses estimating

production functions (and TFP) using Orbis data (Andrews et al., 2016; Gal, 2013), and with Kalemli-Özcan et al. (2019).

A well-known limitation of using Orbis accounts data is that it usually underrepresents smaller and younger firms, depending on the country of reporting. In addition, differences in reporting requirements (for example, in the United States compared to Europe) exist and might also compromise the representativeness and comparability of the data. For this reason, we select listed firms for our analysis. Furthermore, Orbis balance sheet data contain both consolidated (C) and unconsolidated (U) accounts.[26] When a given firm reports both types, we remove duplicates following this order of priority: U1, U2, C1, C2. We drop accounts that report limited financials (which contain rounded or class level figures for a very limited number of variables, particularly affecting countries where detailed financials are not available for most firms, such as the US). Although we follow this strategy to obtain single firm-year observations, some firms in our sample switch account reporting type across years. We assess the consequences of this switching for our baseline DiD estimates in detail in Supplemental Appendix D1.

We then apply some further cleaning steps aimed at removing potential distortions to our estimates arising from erratic or implausible firm performance or behavior. First, we drop the entire firm if the number of persons employed is negative in any year. Second, we assess whether there are unusual data jumps based on growth rates of firms' values of operating revenues. We drop full firm observations if the growth rate is above (below) the top (bottom) percentile.

We interpolate the main output and input variables used in the estimation of the production function if there is one missing value in the middle of the time series for a given firm.

We identify and drop outliers using the ratios of the production function output and inputs.[27] Outliers are defined as any value that is more than three interquartile ranges (IQRs) away from the median. Instead of using the traditional IQR method centered at Q1 and Q3, we use a median-centered (Q2) rule, making it more robust to skewed distributions. Finally, we apply this outlier detection procedure to the level we use for the subsequent markup estimation (i.e., the two-digit NACE sector level). We also apply the same outlier detection method to the markup estimates we computed.

Markups Estimation

The basic principle in De Loecker and Warzynski (2012) is that in a perfectly competitive environment, a firm's cost share for a variable input should equal that input's output elasticity. Deviations from this equality suggest the presence of market power. This technique is applicable across various market structures and does not require specifying a demand system. We assume that each firm i at time t produces output Q_{it} according to a production function that depends

[26] Balance sheet accounts in Orbis are classified according to consolidation status: C1 (consolidated parent and subsidiary account, stand-alone parent account unavailable); C2 (consolidated parent and subsidiary account, stand-alone parent account available); U1 (unconsolidated account, group-level consolidated account unavailable); and U2 (unconsolidated account, group-level consolidated account available).

[27] Namely: capital/labor, output/material costs, output/capital, and output/labor.

on labor L_{it}, capital K_{it}, intermediate inputs M_{it}, and a firm-specific productivity shock Ω_{it}:

$$Q_{it} = F_{it}(L_{it}, K_{it}, M_{it}, \Omega_{it}) \tag{A5}$$

This function is assumed to be continuous and twice differentiable. As in De Loecker and Warzynski (2012), the key identifying assumption is that firms can freely adjust at least one input, in our case, intermediate inputs, which are considered flexible. In contrast, capital and labor are treated as dynamic inputs, reflecting adjustment costs and rigidities.

From the firm's cost minimization problem, we obtain the following markup formula:

$$\mu_{it} = \frac{\theta_{it}^{M}}{\alpha_{it}^{M}} \tag{A6}$$

Here, θ_{it}^{M} denotes the output elasticity of intermediate inputs (to be estimated), and α_{it}^{M} is the share of spending on intermediates relative to sales, which can be directly calculated from balance sheet data. The method's advantage lies in its reliance on just these two quantities.

Estimating θ_{it}^{M} presents a challenge because output elasticities are not directly observable and are potentially biased due to the endogeneity of input choices with respect to productivity. To address this, we use the control function approach of Ackerberg et al. (2015), which effectively corrects for the simultaneity between input decisions and productivity shocks. For structural estimation of production functions, it is common practice to use accounting data, since information on produced quantities is not available. We use balance sheet data on sales (approximating output), total fixed assets (capital, measured by the Orbis variable `fias`, which includes tangible and intangible fixed assets), material expenses (intermediate inputs),[28] and employee costs (labor). All of these variables are converted from nominal to real terms using country-level GDP deflators from the World Bank. A detailed description of the variables included can be found in the following table:

Table A1: Description of variables used in the markup estimation

Variable	Definition and source
Operating revenue (Turnover)	Total operating revenues (i.e., net sales + other operating revenues + stock variations). The figures do not include VAT. Source: Orbis.
Fixed assets (`fias`)	Total amount (after depreciation) of non current assets (i.e., intangible assets + tangible assets + other fixed assets). Source: Orbis.
Cost of employees	Detail of all the employees costs of the company (including pension costs). Source: Orbis.

[28] For observations reporting both material expenses (`mate`) and cost of goods sold (`cost`), we calculate the ratio `mate/cost` and average it within two-digit NACE sectors, pooling across years. When material expenses are missing but `cost` is available, we impute them by multiplying `cost` by this sector-average ratio. The variable `mate` covers purchases of raw materials and finished goods, excluding services; `cost` covers costs of goods sold, production, and services, including depreciation associated with those costs.

Variable	Definition and source
Cost of goods sold	Cost of sold goods, production, and services. Includes costs directly related to the production of the goods sold plus the depreciation of those costs. Source: Orbis.
Material costs	Detail of the purchases of goods (raw materials + finished goods). No services. Source: Orbis.
GDP deflator	Retrieved from the World Bank database. The precise reference is the following: NY.GDP.DEFL.ZS - GDP deflator (base year varies by country).

We estimate production functions separately for each aggregate two-digit NACE sector using a log-linearized translog specification:

$$q_{it} = f_{it}(l_{it}, k_{it}, m_{it}; \beta) + \omega_{it} + \varepsilon_{it} \tag{A7}$$

where the production function $f_{it}(.)$ is assumed to be translog, ω_{it} is the Hicks-neutral productivity level observed by the firm, potentially correlated with the level of inputs in production, while ε_{it} is the unanticipated productivity shock assumed to be independent and identically distributed across firms. Although this approach relies on revenue data rather than physical quantities, which may introduce some bias in levels, prior research such as De Ridder et al. (2026) indicates that this bias is limited when focusing on changes over time, which is our primary interest. The markup is then computed as the ratio of the estimated intermediate input elasticity to its revenue share, correcting for first-stage errors.

Supplemental Appendix B: Patent Data Construction

B1 IP5 Patent Families

We identify the patent offices represented within each family using publication and application country information from Orbis IP. The two fields coincide except when the publication country is coded as "WO", the application country is missing, or the fields use different historical country codes. The code "WO" refers to international applications filed through the World Intellectual Property Organization (WIPO) under the Patent Cooperation Treaty (PCT). We identify the patent office using the publication country, falling back on the application country when the former is missing or coded as "WO". We drop observations for which both the publication and application country are missing. We harmonize historical country codes by reassigning "CS" (Czechoslovakia) to "CZ" (Czech Republic), "YU" (Yugoslavia) to "RS" (Serbia), and "SU" (Soviet Union) to "RU" (Russia). We classify a patent family as IP5 if it includes filings in at least two distinct patent offices worldwide, with at least one filing in the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the United States Patent and Trademark Office (USPTO), or the China National Intellectual Property

Administration (CNIPA).

B2 Digital and AI Patent Classification

We identify digital general-purpose technologies (dGPTs) and artificial intelligence (AI) using a list of tagged patents provided directly by WIPO. The tags follow the classification developed in WIPO (2022) and updated in 2023, which combines International Patent Classification (IPC) and Cooperative Patent Classification (CPC) codes with keyword searches. The underlying patents were extracted in November 2024 from PATSTAT Global. The classification covers AI, autonomous systems, big data, blockchain, cloud computing, digital twins, the Internet of Things (IoT), the metaverse, quantum technologies (computing, entanglement, secure communications, and sensors), robotics, and smart cities.

We link the tagged PATSTAT records to Orbis IP using publication country, publication number, and kind code. When necessary, we use application country, application number, and the corresponding kind code. We then aggregate the tags to the patent-family level. A family is classified as AI if it contains at least one AI-tagged patent. We construct a separate digital indicator covering the remaining dGPT categories, assigning it to families containing at least one patent tagged in any of these categories.

B3 Technological Proximity

We construct each firm's technological profile from the IPC codes associated with its accumulated patent portfolio. For each acquirer-target pair, let $\mathbf{x}$ and $\mathbf{y}$ denote binary vectors whose entries equal one if the corresponding IPC class is represented in the firm's portfolio and zero otherwise. Each class receives equal weight regardless of the number of patents assigned to it. We exclude acquirer-target pairs for which no IPC codes are available for either the acquirer or the target. We measure technological proximity using cosine similarity, defined as $\cos(\theta) = \frac{\mathbf{x}\cdot\mathbf{y}}{\|\mathbf{x}\|\|\mathbf{y}\|}$. As an alternative, we calculate Jaccard similarity, defined as $J(A,B) = \frac{|A\cap B|}{|A\cup B|}$, where A and B are the sets of IPC classes represented in the acquirer's and target's portfolios. Our baseline measure uses four-digit IPC codes; we also construct both measures at one- and three-digit aggregation levels.

For binary IPC vectors, cosine similarity equals the number of shared IPC classes divided by the geometric mean of the numbers of classes represented in the two portfolios, whereas Jaccard similarity divides the same overlap by the number of distinct classes represented in either portfolio. The measures therefore normalize technological overlap differently. For nonempty portfolios, both measures range from zero to one: zero indicates no shared IPC classes, while one indicates identical sets of represented classes. We use cosine similarity in the baseline analysis and Jaccard similarity to assess the sensitivity of the results to this normalization.

For each acquisition, we calculate technological proximity using the acquirer's patent portfolio before incorporating the target's IPC classes. The target's IPC classes enter the acquirer's portfolio in the year following the acquisition. For acquirers completing multiple technological

acquisitions in the same year, we average the deal-level similarity measures across all targets acquired in that year.

Supplemental Appendix C: Descriptive Statistics

Table C1: Number of Firms by Treatment Cohort

Treatment Cohort	Number of firms
2012	132
2013	139
2014	187
2015	190
2016	165
2017	160
2018	164
2019	136
Total treated	1,273
Never treated	6,523
Total	7,796

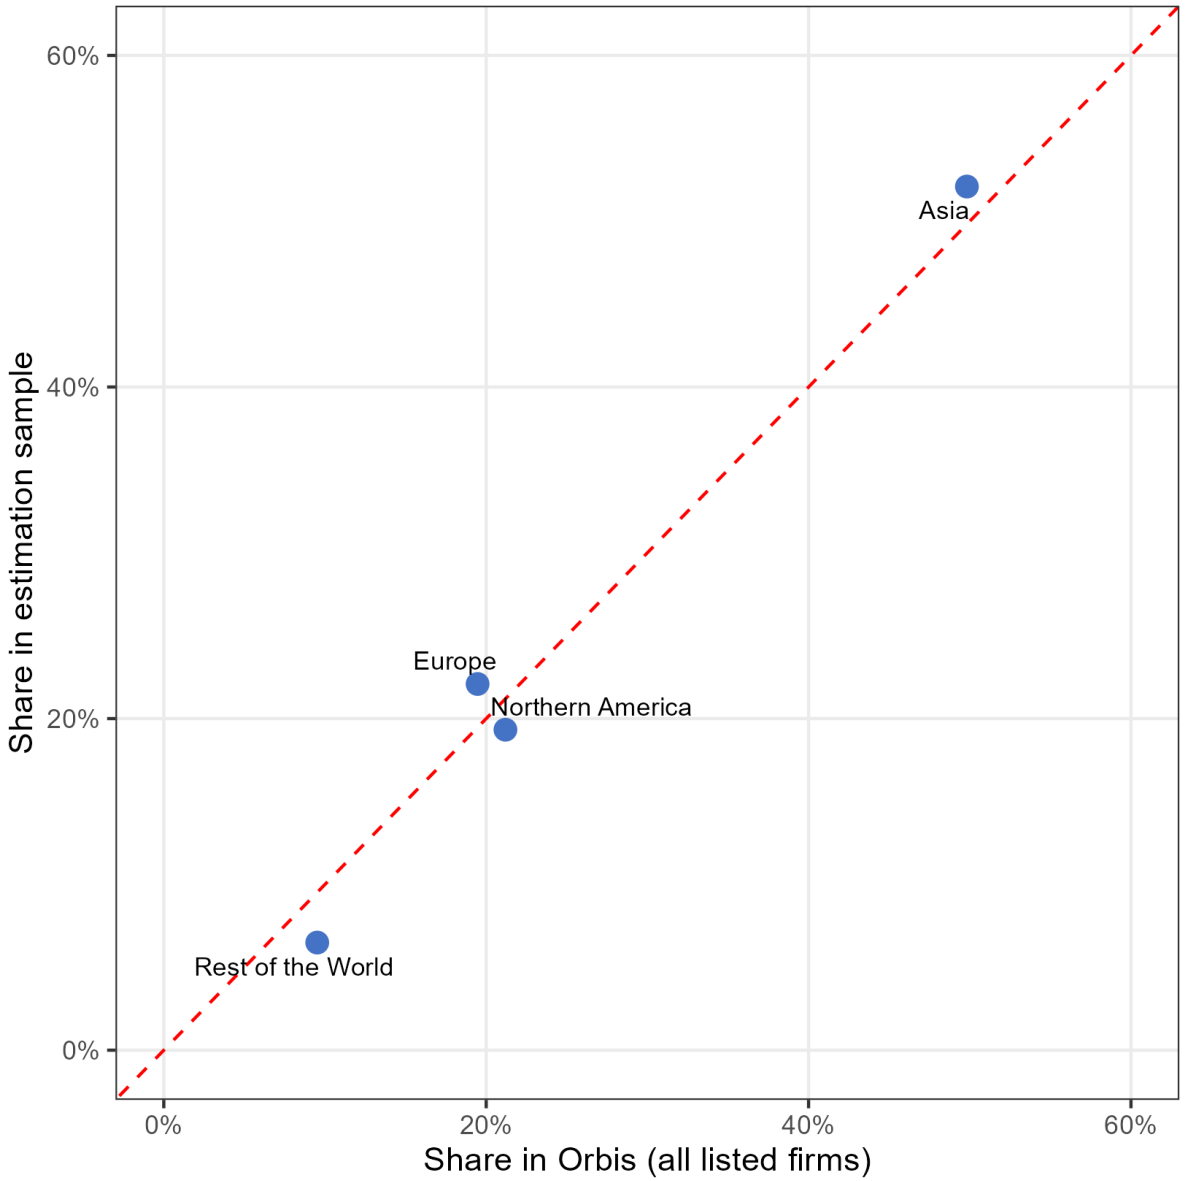
Notes: Number of firms in the estimation sample by year of their first technological M&A. Never treated: acquirers with no technological M&A during the sample period.

Table C2: Sample Distribution by Region

Region	Firms	% firms	Observations	% obs
Asia	4,061	52.1%	38,975	52.9%
Europe	1,722	22.1%	16,464	22.3%
Northern America	1,507	19.3%	13,914	18.9%
Rest of the World	506	6.5%	4,337	5.9%

Notes: Firms and firm-year observations in the estimation sample by region of the country of the firm, following the United Nations M49 regions, with the Americas split into Northern America and Latin America and the Caribbean. Rest of the World comprises Africa, Latin America and the Caribbean, and Oceania.

Figure C1: Sample Representativeness by Region



Notes: The figure compares the share of firms by region in the estimation sample against the share in Orbis (all listed firms). The dashed red line is the 45-degree line. Regions above (below) the line are over- (under-)represented in the estimation sample relative to Orbis.

Table C3: Pre-treatment Balance by Comparison Group

	NY + Pre + Post (baseline)	NY + Post	NY only	Before own first acquisition
Panel A. Levels in the year before the first technological M&A				
Markup	-0.20	-0.18	0.03	-0.25
Real operating revenue (millions)	0.21	0.13	0.11	0.24
Capital stock / total assets	0.02	-0.01	0.12	0.04
Firm age (years)	0.20	0.15	0.01	0.30
Patent stock (PIM)	0.26	0.23	0.12	0.31
Patent applications (families per year)	0.23	0.20	0.09	0.27
Firms holding patents (%)	0.74	0.64	0.11	0.83
Comparison firms	7,595	5,786	1,111	2,698
Panel B. Changes between three years and one year before the first technological M&A				
Change in markup	-0.02	-0.02	0.01	-0.02
Revenue growth (log change)	0.06	0.06	0.02	-0.03
Change in patent applications (IHS)	0.09	0.09	0.04	0.07
Comparison firms	7,058	5,471	1,022	2,166
Treated firms	1,091 (Panel A); 1,042 (Panel B)			

Notes: Normalized differences between technological acquirers and comparison firms, computed as in Table 1, for the comparison groups of Table 2, restricted to acquisition cohorts 2012–2018 (the 2019 cohort has no later-treated comparison firms). NY: firms whose first technological acquisition occurs after the focal cohort year; Pre and Post: acquirers with no recorded technological acquisition, observed before and after their own first acquisition. The last column compares technological acquirers with acquirers with no recorded technological acquisition whose first acquisition occurs in the focal cohort year, both observed in the year before. Comparison firms are distinct firms contributing to the markup row; treated firms are those with the characteristic available.

Supplemental Appendix D: Robustness Checks

In this section, we present the robustness of our difference-in-differences (DiD) estimates to: (1) account consolidation switching within firms, (2) the markup measure, (3) the overall identification strategy (the use of covariates and an alternative comparison group, alternative non-absorbing treatment definition for Section IIIA), and (4) the similarity measures used.

D1 Account Consolidation

Despite removing duplicate account types in the established order described in Supplemental Appendix A2, there is a non-negligible number of firms that switch account reporting type across years (33.4 percent of firms), so that the panel may blend balance sheet definitions for the same firm over time.

Such changes are a potential threat to the reliability of both the markup measure and the identification strategy. First, consolidating a subsidiary into the group accounts mechanically adds the subsidiary's external revenue and eliminates the inter-group transactions between parent and subsidiary (Levonyan and Mengano, 2024). These two adjustments scale revenue and materials expenditure by different amounts, so that the materials expenditure share that enters the markup formula can shift even when the firm's underlying market power stays the same. Second, the bias is compounded in our setting because a switch from unconsolidated

to consolidated reporting is most likely to occur precisely when a firm conducts a technological acquisition: were treated firms systematically more likely to switch toward consolidated reporting in the post-treatment period, this would generate a mechanical, treatment-correlated upward bias in our estimated markup effect.

We therefore conduct diagnostic checks for potential mixing of unconsolidated and consolidated accounts in our sample. We find no evidence that this mechanism is at play: conditional on firm and industry-by-year fixed effects, treated firms become, if anything, less likely to be on consolidated accounts in the post-treatment period (by 6.5 percentage points in a linear probability model), and the within-firm markup gap between consolidated and unconsolidated observations is statistically indistinguishable from zero (Column (1) of Table D1).

Table D1: DiD Robustness: Restricting to Consistent Consolidation Type

	(1) Cons. ctrl	(2) Constant	(3) Switcher	(4) Unconsol. only	(5) Consol. only
M&A tech	0.0179***	0.0289***	-0.0021	0.0106*	0.0294***
	(0.0044)	(0.0058)	(0.0064)	(0.0060)	(0.0061)
Consolidated accounts	0.0016				
	(0.0035)				
Firm FE	Yes	Yes	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes	Yes	Yes
Acct type control	Yes	No	No	No	No
R sq	0.835	0.847	0.817	0.853	0.856
N (firm X year)	73690	48110	25565	34821	38449

Notes: Robustness of the main DiD markup estimates when controlling for the presence of consolidated accounts (Column (1)) and restricting the sample to firms that report different consolidation types: (i) no within-firm switches between unconsolidated and consolidated accounts (Column (2)), (ii) within-firm switches between unconsolidated and consolidated accounts (Column (3)), (iii) firms with unconsolidated accounts only (Column (4)), and (iv) firms with consolidated accounts only (Column (5)). All estimations include firm and industry-by-year fixed effects, with standard errors (displayed in parentheses) clustered by firm. $^{*}\ p < 0.1$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.

D2 Profitability and Excess Markup

If the estimated markup increase reflects a genuine gain in market power rather than a mechanical shift in the cost structure due to reinvestment of profits in infrastructure, organizational capital, or other intangible and fixed assets, it should also be visible in firms' realized profitability and should survive the removal of fixed operating costs. To this end, we report two complementary checks.

First, Table D2 reports the effect of a technological M&A on three profit measures: the operating margin (operating profit over revenue), operating profit, and operating revenue, the latter two inverse-hyperbolic-sine transformed. In the TWFE specification all three increase (0.0065, 0.109, and 0.095, respectively). The Callaway and Sant'Anna (2021) estimates (Table D3)

tell the same story, with a positive and significant effect on operating profit (0.029) and revenue (0.051), with a positive but less precisely estimated effect on the operating margin (0.005).

As an additional check, Table D4 reports the effect on excess markup. This is the Lerner index net of the industry-by-year benchmark fixed-cost share, as defined in Section IIB1. This follows De Loecker et al. (2020) and De Ridder (2024): under the identifying assumption that long-run economic profits are zero within an industry-by-year cell, the average Lerner index in that cell should just cover the average firm's fixed costs, so the industry-by-year average price-cost margin proxies the fixed-cost share a firm needs to break even. Excess markup therefore nets out the portion of the margin that is plausibly absorbed by ordinary fixed-cost recovery (including the reinvestment channels noted above) and isolates the residual margin that is not justified by fixed-cost coverage, i.e., a closer proxy for pure market power rents. The effect is positive and significant under both estimators (0.0066 with TWFE and 0.0098 with Callaway and Sant'Anna (2021)), consistent with the markup gains surviving the removal of the fixed-cost component and with market power, rather than a mechanical shift in the cost structure, driving the increase.

Table D2: Effect of Technological M&A on Profitability (TWFE)

	(1) Prof. marg.	(2) OP	(3) OR
M&A tech	0.0065**	0.1093***	0.0950***
	(0.0028)	(0.0420)	(0.0185)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
F-test	5.423	6.764	26.370
R sq	0.508	0.784	0.951
N (firm X year)	71230	71299	76807

Notes: TWFE estimates of the effect of a technological M&A on profitability. The dependent variables are the operating margin (operating profit over operating revenue), operating profit (OP), and operating revenue (OR); OP and OR are inverse-hyperbolic-sine transformed (in millions). All specifications include firm and industry-by-year fixed effects. Standard errors in parentheses, clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D3: Effect of Technological M&A on Profitability (Callaway & Sant'Anna)

	(1) Prof. marg.	(2) OP	(3) OR
M&A tech	0.0052	0.0294**	0.0511***
	(0.0041)	(0.0115)	(0.0187)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
N (firm X year)	71010	71079	76676

Notes: Callaway and Sant'Anna (2021) estimates of the effect of a technological M&A on profitability, with the same outcomes as Table D2. Outcomes are pre-residualized on industry-by-year means among never-treated firms. Standard errors in parentheses, clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table D4: Effect of Technological M&A on Excess Markup

	(1)	(2)
M&A tech	0.0066**	
	(0.0031)	
M&A tech		0.0098***
		(0.0026)
Estimator	TWFE	CS
Firm FE	Yes	Yes
Ind X Year FE	Yes	Yes
F-test	4.581	
R sq	0.823	
N (firm X year)	68277	67997

Notes: Effect of a technological M&A on excess markup, defined as the Lerner index $(1 - 1/\mu)$ net of the industry-by-year benchmark fixed-cost share (the leave-one-out, sales-weighted average fixed-cost share in the firm's NACE-by-year cell). Column (1) reports TWFE estimates; Column (2) reports Callaway and Sant'Anna (2021) estimates. Standard errors in parentheses, clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

D3 Covariates and Alternative Comparison Group

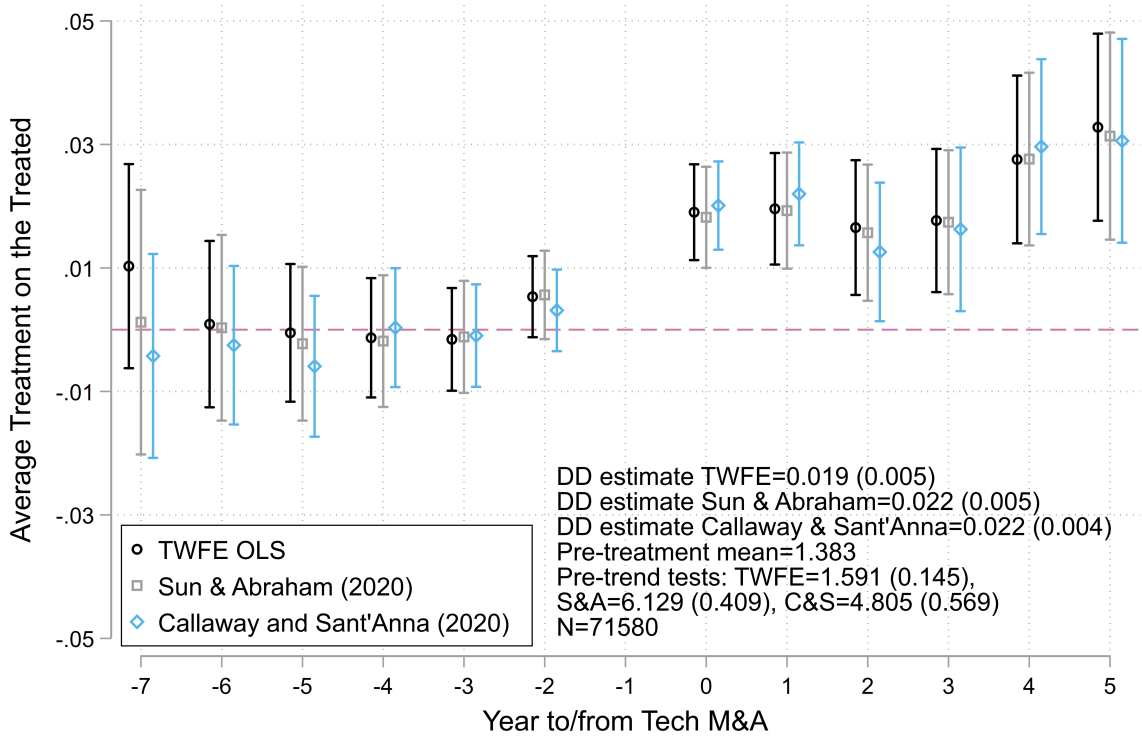
All the DiD estimates presented in the main text in Section III only include firm and industry-by-year fixed effects (or, for the Callaway and Sant'Anna (2021) estimator, outcomes pre-residualized on industry-by-year means). No covariates are included. This is deliberate, following recent econometric work. In a DiD design, covariates are meant to justify a conditional parallel-trends assumption, not to block confounding as in a cross-sectional setting. Including covariates in a standard two-way fixed effects regression can be non-neutral: Caetano and Callaway (2024) show that the within-unit transformation removing the firm fixed effect also changes the covariates (conditioning on growth rather than level), inducing a hidden linearity bias. An additional issue is that covariates measured post-treatment are themselves potentially affected by treatment (Baker et al., 2026). We therefore chose not to include covariates in our main specification, also because pre-treatment trends are parallel unconditionally in our case. The same considerations apply to matching, which imposes parallel trends conditional on the matching variables; we therefore use it as a robustness check rather than in the baseline, also because one-to-one matching drops treated firms without a match and changes the estimand to the matched subsample.

Nevertheless, we check the robustness of the baseline event-study in Section IIIA along two dimensions: the inclusion of pre-treatment covariates and the use of a matched comparison group. As for the former, we restrict any covariates used as a robustness check to pre-treatment values. We add log capital stock, log firm age, and IHS-transformed patent stock computed by the perpetual inventory method with 5 percent depreciation, averaged over all years prior to the first technological M&A for treated firms and over the years prior to their first acquisition for firms with no technological M&A (over the full panel when no such year is observed). When using the Callaway and Sant'Anna (2021) estimator, we use the doubly robust estimator of Sant'Anna and Zhao (2020), which has been shown to be less affected by potential bias due to the inclusion of covariates. In the doubly robust estimator, we additionally include log real operating revenue, averaged over the same years. As for the latter, the matched sample is constructed using Coarsened Exact Matching (CEM; Iacus et al. 2012). Matching is performed on two variables, both matched exactly: (i) the NACE industry sector code (two-digit NACE rev. 2), and (ii) the M&A year (year of first technological M&A for treated firms; year of first M&A of any type for untreated comparison firms). We retain one comparison firm per treated firm, so treated firms without a match are dropped.

Figure D1 re-estimates the baseline specification in Figure 2 adding the above-mentioned controls. The average post-treatment estimates are 0.020–0.022 across estimators, essentially identical to the baseline, and the pre-treatment coefficients remain flat. Figures D2 and D3 replace the comparison group with one constructed by Coarsened Exact Matching (CEM). The estimated effect is again stable in magnitude and significance, 0.018–0.023 without controls and 0.021 with controls, so the baseline result in Figure 2 is not driven by the composition of the comparison group. Figure D4 re-estimates the event study using only not-yet-treated firms as the comparison group, the dynamic counterpart of Column (3) of Table 2. The pre-treatment coefficients are jointly insignificant ($p = 0.354$ for TWFE and $p = 0.256$ for Callaway

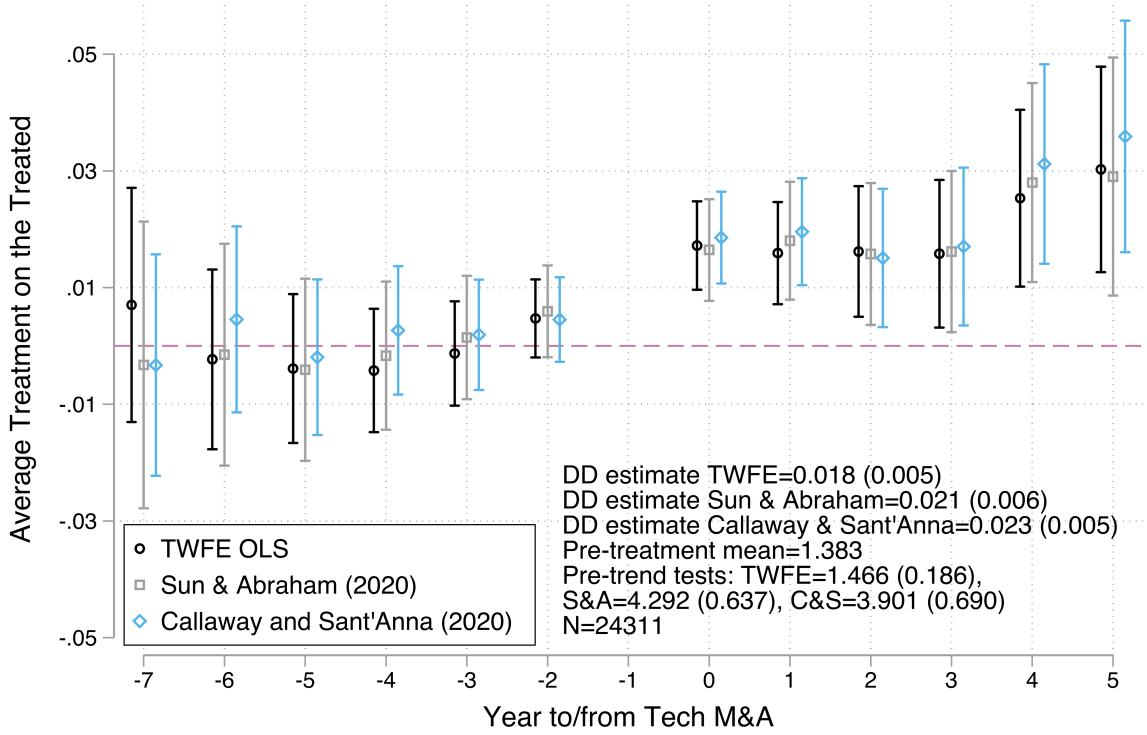
and Sant'Anna (2021)), and the average post-treatment effects, 0.015 and 0.021, are close to the baseline estimates; individual coefficients at longer horizons are less precisely estimated, reflecting the smaller number of not-yet-treated comparison firms observed far from treatment. We do not report the Sun and Abraham (2021) estimator for this sample: with only not-yet-treated firms, its comparison group reduces to the 2019 cohort, whose pre-acquisition markup trend differs from that of the other cohorts. In this comparison, pre-treatment characteristics are also close to balanced, with normalized differences of at most 0.12 in absolute value (Table C3).

Figure D1: Effect of Technological M&A on Markup (with Controls)



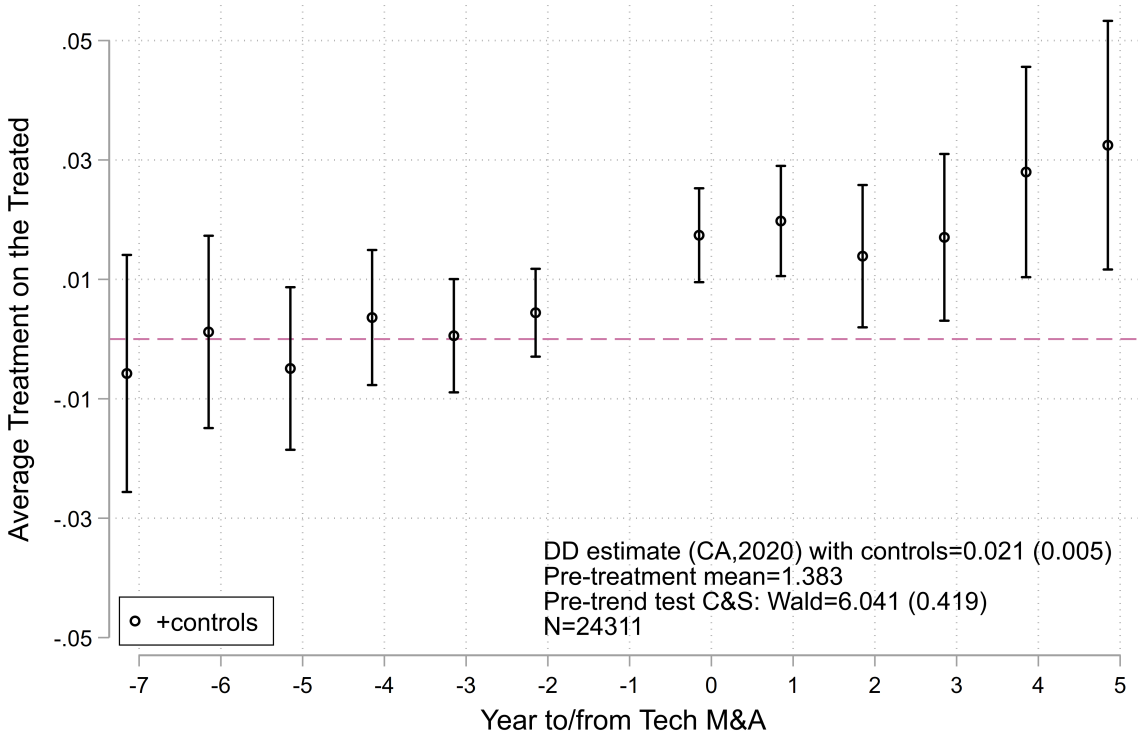
Notes: The figure shows the estimated effects of a technological M&A on firm markup. The figure reports the point estimates and confidence intervals from (i) the two-way fixed effects (TWFE) OLS estimator, (ii) Sun and Abraham (2021), and (iii) Callaway and Sant'Anna (2021). Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal. All effects are relative to the year before the deal. Controls include log-transformed firm age and capital stock, and IHS-transformed patent stock. The TWFE and Sun and Abraham (2021) specifications include firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A, and the estimates also control for log real operating revenue. Callaway and Sant'Anna (2021) estimates adjust for the inclusion of pre-treatment covariate values using the doubly robust methodology discussed in Sant'Anna and Zhao (2020). Confidence intervals are at the 95 percent level and based on standard errors clustered at the firm level.

Figure D2: Effect of Technological M&A on Markup (Matched Sample)



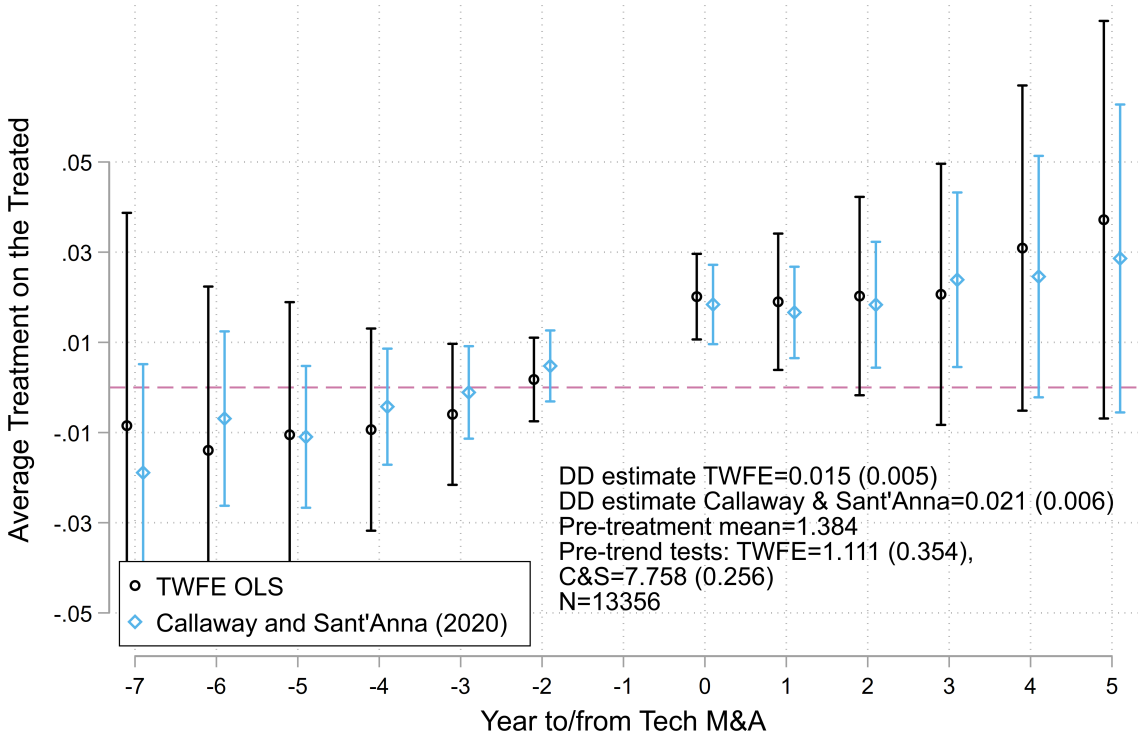
Notes: The figure shows the estimated effects of a technological M&A on firm markup using the Coarsened Exact Matching (CEM) sample. The figure reports the point estimates and confidence intervals from (i) the two-way fixed effects (TWFE) OLS estimator, (ii) Sun and Abraham (2021), and (iii) Callaway and Sant'Anna (2021). Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal. All effects are relative to the year before the deal. The comparison group is matched on two variables, both matched exactly: the two-digit NACE rev. 2 industry sector code and the M&A year (year of first technological M&A for treated firms; year of first M&A of any type for untreated comparison firms), retaining one comparison firm per treated firm. Confidence intervals are at the 95 percent level and based on standard errors clustered at the firm level.

Figure D3: Effect of Technological M&A on Markup (Matched Sample with Controls)



Notes: The figure shows the estimated effects of a technological M&A on firm markup using the Callaway and Sant'Anna (2021) estimator on the Coarsened Exact Matching (CEM) sample, adding pre-treatment covariate controls. Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal. All effects are relative to the year before the deal. The comparison group is matched on the two-digit NACE rev. 2 industry sector code and the M&A year, retaining one comparison firm per treated firm. Controls are pre-treatment firm-level means of log capital stock, log firm age, and IHS-transformed patent stock via the perpetual inventory method with 5 percent depreciation, adjusted for using the doubly robust methodology discussed in Sant'Anna and Zhao (2020). Confidence intervals are at the 95 percent level and based on standard errors clustered at the firm level.

Figure D4: Effect of Technological M&A on Markup (Not-Yet-Treated Controls Only)



Notes: The figure shows the estimated effects of a technological M&A on firm markup when the comparison group comprises only firms whose first technological M&A occurs later (not-yet-treated firms), as in Column (3) of Table 2. The figure reports the point estimates and confidence intervals from (i) the two-way fixed effects (TWFE) OLS estimator and (ii) Callaway and Sant'Anna (2021). Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal. All effects are relative to the year before the deal. The TWFE specification includes firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A. Confidence intervals are at the 95 percent level and based on standard errors clustered at the firm level.

D4 Non-Absorbing Treatment Definition (Robustness to Reversals)

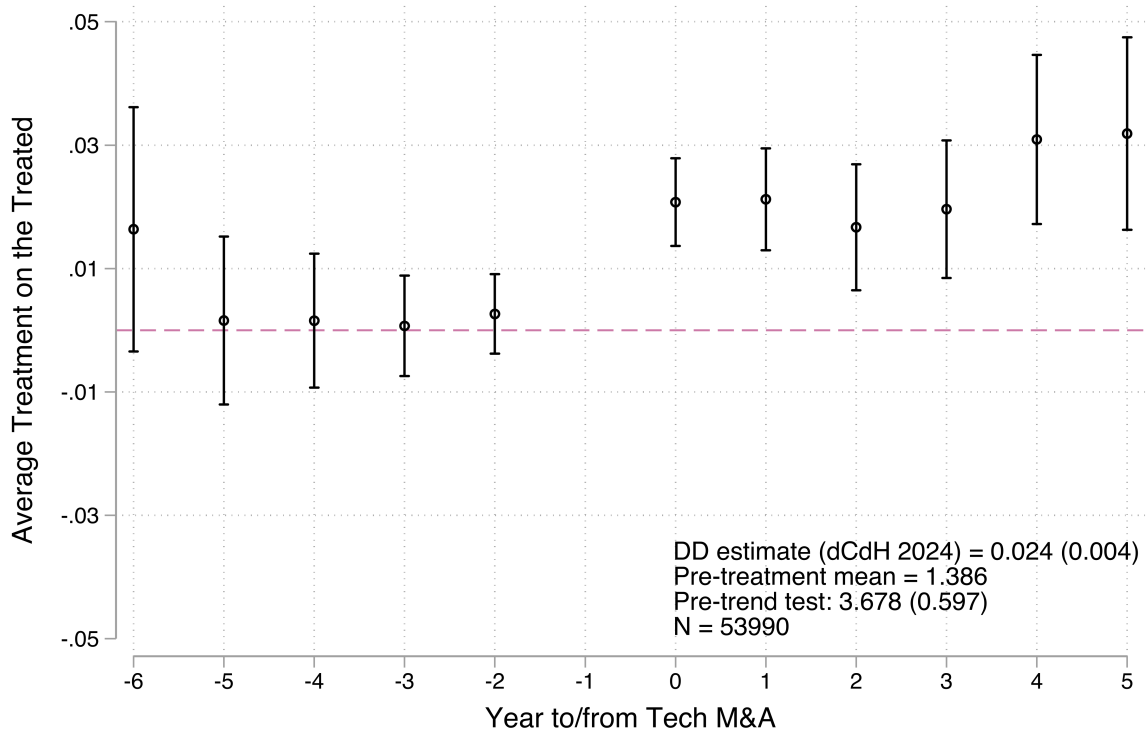
The treatment, conducting a technological M&A, is non-absorbing by nature: the data record each technological M&A as a discrete, year-specific deal. However, in the baseline analysis in Sections IIA (description of empirical strategy) and IIIA (results) we convert it into an absorbing treatment, equal to one from the year of a firm's first technological M&A onward, on the grounds that the capabilities and market position acquired through the deal are unlikely to be undone.[29] In this section we conduct a robustness check on whether changing the treatment definition might bias our baseline results, by re-specifying a non-absorbing treatment (dummy equal to one only in the year of each technological deal). This definition implies that some firms switch into and out of treatment over time.

Standard heterogeneity-robust DiD estimators, including Callaway and Sant'Anna (2021) and Sun and Abraham (2021), are built for binary and staggered (absorbing) treatments, and applying them to a treatment that can switch on and off can violate the assumptions under which they are valid (de Chaisemartin and D'Haultfœuille, 2020). To check whether the non-absorbing treatment definition drives our results in Section IIIA (Figure 2), we re-estimate the baseline markup regression using the heterogeneity-robust event-study estimator of De Chaisemartin and d'Haultfoeuille (2024). This estimator accommodates treatments that are non-binary, non-absorbing, and that may switch in and out multiple times. For each switcher (a firm whose treatment changes during the panel) it computes a sequence of event-study effects comparing the firm's markup evolution, from the period before its treatment first changes onward, to that of firms sharing the same baseline treatment status whose treatment has not yet changed. It then aggregates only across these pairs, thus returning an estimator that is robust to potential bias from switching status.

The average post-treatment effect and the joint test on the five placebo coefficients, reported in Figure D5, are in line with those obtained under the absorbing treatment definition used throughout the rest of the paper, confirming that treatment reversals do not drive the main findings.

[29] Recall that the mechanism tests in Section IIIC that exploit post-acquisition patent flow and patent stock instead revert to a non-absorbing treatment—a dummy equal to one only in the year of each technological deal—because these tests are designed to capture the immediate post-acquisition dynamics of a specific transaction.

Figure D5: Effect of Technological M&A on Markup, Robust to Reversals (De Chaisemartin and d'Haultfoeuille, 2024)

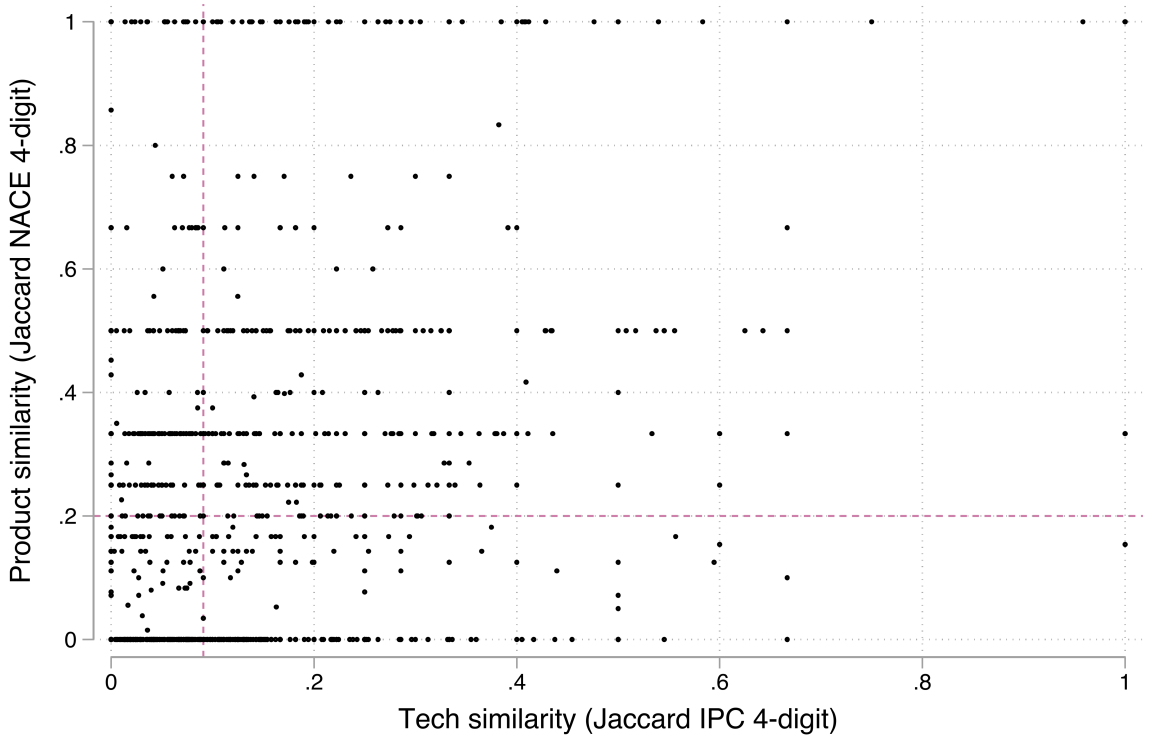


Notes: The figure shows the estimated effects of a technological M&A on firm markup under the non-absorbing treatment definition. The figure reports the coefficients and standard errors from the heterogeneity-robust event-study estimator of De Chaisemartin and d'Haultfoeuille (2024), which compares switchers—firms whose treatment status changes during the panel—to firms sharing the same baseline treatment that have not yet switched. Coefficients are shown for up to six years before the deal, the deal year, and up to five years after the deal. All effects are relative to the year before the deal. Markups are pre-residualized on industry-by-year means of firms with no technological M&A. Confidence intervals are at the 95 percent level and based on standard errors clustered at the firm level.

D5 Similarity Analysis with Jaccard Measure

Figure D6 reproduces Figure 1 (Section IIB2) using Jaccard instead of cosine similarity.

Figure D6: Similarity Scatter Plot (Jaccard)



Notes: Each point represents an acquirer. The axes show the maximum observed annual Jaccard similarity in technological space (four-digit IPC classes, x-axis) and product-market space (four-digit NACE sectors, y-axis), calculated separately across its acquisitions. The dashed lines indicate the median of each similarity measure computed among ever-treated firms, defining the four quadrants used in the regression analysis: HH (top right), HL (bottom right), LH (top left), and LL (bottom left).

Table D5 replicates Table 7 (Section IIIC2) using two quadrant dummies based on Jaccard rather than cosine similarity; the results confirm those in the main text.

Table D5: Technological M&A and Proximity: HH vs Rest (Jaccard Similarity)

	(1) TWFE	(2) CS: HH	(3) CS: Rest
M&A tech × HH	0.0373***		
	(0.0091)		
M&A tech × Rest (HL+LH+LL)	0.0110**		
	(0.0054)		
M&A tech		0.0487***	0.0181***
		(0.0083)	(0.0052)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
F: HH=Rest	6.465		
p: HH=Rest (TWFE)	0.011		
z: HH=Rest		3.130	
p: HH=Rest (CS z-test)		0.002	
F-test (model)	10.172		
R sq	0.836		
N (firm X year)	73358	63356	67522

Notes: Robustness using Jaccard similarity in place of cosine similarity. TWFE and Callaway and Sant'Anna (2021) estimates comparing HH firms (high tech similarity + high product-market similarity) against the rest (HL+LH+LL). Technology similarity is the Jaccard index of four-digit IPC patent classes; product-market similarity is the Jaccard index of four-digit NACE sector codes. Median split among ever-treated firms. The test statistic and p-value for the equality of HH and Rest coefficients are reported. The TWFE specification includes firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors in parentheses are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

D6 Similarity Analysis with Cosine Measure (Four Quadrants)

The discussion in Section IIB2 and results presented in Table 7 in Section IIIC2 focus on comparing the HH quadrant (i.e., transactions where both technological and market proximity are high) with the rest of transactions in the remaining quadrants. The HL quadrant (high technological and low market proximity; bottom right in Figure 1) is also of interest. For these transactions, patent densification (i.e., the acquired patents densify the acquirer's portfolio around its core technological domain) remains a plausible motive, but the absence of direct product-market rivalry makes the interpretation less clear-cut: high technological proximity also facilitates the evaluation and integration of the target's knowledge base through absorptive capacity (Cassiman et al., 2005; Cohen et al., 1990), leaving open a role for innovation synergies alongside shielding. We expect a positive markup effect in this quadrant, but potentially attenuated relative to HH and with a more ambiguous mechanism. When technological proximity is low—regardless of market proximity (LH and LL quadrants)—the scope for patent portfolio densification is limited, and we expect weaker and less systematic markup effects.

If the innovation-shielding channel is the dominant mechanism, markup gains should be largest and most precisely estimated in the HH quadrant, and should attenuate as technological proximity falls. Comparing HH and HL allows us to assess the marginal contribution of market proximity—and hence direct rivalry elimination—beyond the patent densification motive captured by technological proximity alone.

Table D6 presents more detailed results using four quadrant dummies separately (i.e., instead of only HH vs the rest), using the cosine similarity measure. They also confirm the results presented in Table 7 in Section IIIC2, while providing further nuances. Statistically significant markup effects are concentrated among deals with above-median technological similarity, consistent with the innovation-shielding channel being central to the markup effect of technological M&As. The strongest effect is found in the HH quadrant (CS ATT of 0.050, against 0.024 in the baseline), consistent with the innovation-shielding channel. The HL quadrant also yields a positive and significant effect (0.034), consistent with innovation synergies as a secondary channel (Cassiman et al., 2005). The two coefficients are not statistically distinguishable, likely reflecting limited power from the four-way partition of the treated group.

Deals with low technological similarity (LH and LL) yield small and largely insignificant effects regardless of product-market overlap. Moreover, the HH coefficient differs statistically from both LH and LL, suggesting that technological proximity is the dimension along which the markup effect concentrates, with product-market proximity playing a secondary role.

Table D6: Technological M&A and Proximity: Four Quadrants (Cosine Similarity)

	(1) TWFE	(2) CS: HH	(3) CS: HL	(4) CS: LH	(5) CS: LL
M&A tech × HH	0.0389***				
	(0.0091)				
M&A tech × HL	0.0272**				
	(0.0108)				
M&A tech × LH	0.0010				
	(0.0091)				
M&A tech × LL	0.0061				
	(0.0079)				
M&A tech		0.0502***	0.0341***	0.0164*	0.0058
		(0.0081)	(0.0094)	(0.0093)	(0.0078)
Firm FE	Yes	Yes	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes	Yes	Yes
F: HH=HL	0.701				
p: HH=HL	0.403				
F: HH=LH	8.866				
p: HH=LH	0.003				
F: HH=LL	7.588				
p: HH=LL	0.006				
F: HL=LH	3.517				
p: HL=LH	0.061				
F: HL=LL	2.558				
p: HL=LL	0.110				
F: LH=LL	0.184				
p: LH=LL	0.668				
z: HH=HL		1.292			
p: HH=HL		0.197			
z: HH=LH		2.732			
p: HH=LH		0.006			
z: HH=LL		3.946			
p: HH=LL		0.000			
z: HL=LH		1.343			
p: HL=LH		0.179			
z: HL=LL		2.331			
p: HL=LL		0.020			
z: LH=LL		0.880			
p: LH=LL		0.379			
F-test (model)	6.084				
R sq	0.836				
N (firm X year)	73358	63420	62224	62475	63279

Notes: TWFE and Callaway and Sant'Anna (2021) estimates by similarity quadrant (cosine four-digit IPC and NACE). HH: high tech + high product similarity; HL: high tech + low product; LH: low tech + high product; LL: low tech + low product. Pairwise test p-values reported. The TWFE specification includes firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors in parentheses are clustered at the firm level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Supplemental Appendix E: Regression Tables for Figures

This section reports in tabular form the estimates plotted in the event-study and heterogeneity figures of the main text and of the Supplemental Appendix. Each table gives the full set of coefficients, standard errors, and summary statistics behind the corresponding figure.

Table E1: Effect of Technological M&A on Markup (Figure 2)

	(1) TWFE OLS	(2) Sun & Abraham	(3) Callaway & Sant'Anna
Year -7	0.0104	0.0013	-0.0002
	(0.0084)	(0.0108)	(0.0086)
Year -6	0.0016	-0.0002	0.0038
	(0.0070)	(0.0077)	(0.0068)
Year -5	0.0005	-0.0025	-0.0024
	(0.0057)	(0.0064)	(0.0058)
Year -4	-0.0017	-0.0026	0.0008
	(0.0049)	(0.0054)	(0.0049)
Year -3	-0.0004	0.0000	0.0003
	(0.0043)	(0.0047)	(0.0043)
Year -2	0.0050	0.0053	0.0029
	(0.0033)	(0.0036)	(0.0033)
Year 0 (deal)	0.0190***	0.0183***	0.0208***
	(0.0038)	(0.0040)	(0.0036)
Year +1	0.0179***	0.0176***	0.0212***
	(0.0043)	(0.0046)	(0.0042)
Year +2	0.0153***	0.0144***	0.0167***
	(0.0054)	(0.0054)	(0.0052)
Year +3	0.0154***	0.0146**	0.0196***
	(0.0058)	(0.0059)	(0.0057)
Year +4	0.0251***	0.0249***	0.0309***
	(0.0069)	(0.0071)	(0.0070)
Year +5	0.0304***	0.0293***	0.0319***
	(0.0077)	(0.0085)	(0.0080)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
Average post-treatment ATT	0.0176	0.0199	0.0238
(s.e.)	0.0045	0.0046	0.0041
Pre-trend test (p-value)	0.253	0.506	0.803
Pre-treatment mean markup	1.386	1.386	1.386
N (firm X year)	73690	73690	73690

Notes: Tabular version of Figure 2. The table reports the estimated effects of a technological M&A on firm markup from: (i) the two-way fixed effects (TWFE) OLS estimator, (ii) Sun and Abraham (2021), and (iii) Callaway and Sant'Anna (2021). Coefficients are shown for up to seven years before the first technological M&A, the deal year, and up to five years after the deal; all effects are relative to the year before the deal. The TWFE and Sun and Abraham (2021) specifications include firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E2: Effect of Technological M&A on Markup, Top R&D Investors (Figure 3a)

	(1) Top R&D investors	(2) Other firms
M&A tech	0.0360*** (0.0079)	0.0116** (0.0053)
Firm FE	Yes	Yes
Ind X Year FE	Yes	Yes
p-value: equality across columns	0.010	
N (firm X year)	11346	62094

Notes: Tabular version of Figure 3a. The table reports the average post-treatment effect of a technological M&A on firm markup from the estimator of Callaway and Sant'Anna (2021), estimated separately for top R&D investors and for all other firms. Top R&D investors are firms that appeared at least once in the list of the top 2,500 R&D investors during the period 2003–2024 (Nindl et al., 2024). Markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm. The reported p-value is for the null of equal effects across the two columns, based on an approximate z-test.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E3: Effect of Technological M&A on Markups across Regions (Figure 3b)

	(1) China	(2) EU	(3) Japan	(4) RoW	(5) US
M&A tech	0.0041 (0.0128)	0.0519*** (0.0165)	0.0084 (0.0065)	0.0123 (0.0103)	0.0522*** (0.0080)
Firm FE	Yes	Yes	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes	Yes	Yes
N (firm X year)	15188	9283	13792	22370	12789

Notes: Tabular version of Figure 3b. The table reports the average post-treatment effect of a technological M&A on firm markup from the estimator of Callaway and Sant'Anna (2021), estimated separately by world region of the acquirer. RoW denotes the rest of the world. Markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E4: Effect of Technological M&A on Markups across Industries (Figure 3c)

	(1) HT manuf	(2) MHT manuf	(3) MLT manuf	(4) LT manuf	(5) KIS	(6) LKIS
M&A tech	0.0388***	0.0175***	0.0045	0.0193*	0.0428***	-0.0071
	(0.0104)	(0.0059)	(0.0088)	(0.0108)	(0.0130)	(0.0102)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N (firm X year)	9214	11911	6514	8580	16164	14026

Notes: Tabular version of Figure 3c. The table reports the average post-treatment effect of a technological M&A on firm markup from the estimator of Callaway and Sant'Anna (2021), estimated separately by major industry aggregate. Industry aggregation follows the Eurostat sectoral approach (Eurostat, 2014): HT, MHT, MLT, and LT manuf. denote high-, medium-high-, medium-low-, and low-technology manufacturing; KIS and LKIS denote knowledge-intensive and less knowledge-intensive services. Markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E5: Heterogeneity by Knowledge-Intensive Services (Figure 3d)

	(1) Market	(2) High tech	(3) Financial	(4) Other
M&A tech	0.0429	0.0499**	-0.0747	0.0544***
	(0.0292)	(0.0199)	(0.0496)	(0.0209)
Firm FE	Yes	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes	Yes
N (firm X year)	3928	7195	1126	3915

Notes: Tabular version of Figure 3d. The table reports the average post-treatment effect of a technological M&A on firm markup from the estimator of Callaway and Sant'Anna (2021), estimated separately within the four knowledge-intensive service categories of the Eurostat sectoral approach (Eurostat, 2014). Markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E6: Effect of IP5, dGPT and AI Technological M&As on Markup (Figure 4)

	(1) IP5 M&A	(2) Digital M&A	(3) AI M&A
Year -7	0.0032	-0.0175	-0.0412*
	(0.0099)	(0.0203)	(0.0243)
Year -6	0.0076	-0.0116	-0.0173
	(0.0078)	(0.0146)	(0.0215)
Year -5	0.0014	-0.0055	-0.0083
	(0.0062)	(0.0124)	(0.0179)
Year -4	0.0042	-0.0002	0.0095
	(0.0052)	(0.0105)	(0.0149)
Year -3	0.0027	0.0090	0.0280**
	(0.0046)	(0.0094)	(0.0122)
Year -2	0.0039	0.0137**	0.0203***
	(0.0034)	(0.0060)	(0.0079)
Year 0 (deal)	0.0225***	0.0561***	0.0785***
	(0.0039)	(0.0102)	(0.0181)
Year +1	0.0208***	0.0498***	0.0646***
	(0.0047)	(0.0127)	(0.0215)
Year +2	0.0165***	0.0363***	0.0578***
	(0.0056)	(0.0134)	(0.0212)
Year +3	0.0226***	0.0428***	0.0589**
	(0.0063)	(0.0152)	(0.0236)
Year +4	0.0309***	0.0634***	0.0812**
	(0.0075)	(0.0209)	(0.0328)
Year +5	0.0331***	0.0504**	0.0852**
	(0.0085)	(0.0253)	(0.0414)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
Average post-treatment ATT	0.0251	0.0494	0.0728
(s.e.)	0.0046	0.0112	0.0183
Pre-trend test (p-value)	0.837	0.359	0.029
Pre-treatment mean markup	1.408	1.442	1.440
N (firm X year)	72590	69137	68768

Notes: Tabular version of Figure 4. The table reports the estimated effects of a technological M&A on firm markup from the estimator of Callaway and Sant'Anna (2021), distinguishing between acquisitions involving only IP5, dGPT, and AI patents, respectively. dGPT and AI patents are identified following the procedure in WIPO (2022). Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal; all effects are relative to the year before the deal. Each column uses its own treatment-specific comparison group. Markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E7: Effect of Technological M&A on Productivity (Figure 5)

	(1) TWFE OLS	(2) Sun & Abraham	(3) Callaway & Sant'Anna
Year -7	-0.0005	-0.0022	0.0043
	(0.0110)	(0.0142)	(0.0125)
Year -6	-0.0002	0.0042	-0.0024
	(0.0088)	(0.0101)	(0.0094)
Year -5	-0.0056	-0.0013	-0.0078
	(0.0076)	(0.0086)	(0.0079)
Year -4	0.0032	0.0082	0.0025
	(0.0065)	(0.0071)	(0.0066)
Year -3	0.0034	0.0053	0.0015
	(0.0057)	(0.0062)	(0.0058)
Year -2	-0.0016	-0.0010	-0.0013
	(0.0047)	(0.0051)	(0.0048)
Year 0 (deal)	-0.0246***	-0.0218***	-0.0256***
	(0.0053)	(0.0056)	(0.0053)
Year +1	-0.0195***	-0.0150***	-0.0221***
	(0.0055)	(0.0057)	(0.0056)
Year +2	-0.0149**	-0.0127**	-0.0149**
	(0.0063)	(0.0064)	(0.0064)
Year +3	-0.0135*	-0.0120	-0.0176**
	(0.0074)	(0.0077)	(0.0078)
Year +4	-0.0072	-0.0013	-0.0097
	(0.0080)	(0.0084)	(0.0088)
Year +5	-0.0061	-0.0054	-0.0133
	(0.0099)	(0.0113)	(0.0110)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
Average post-treatment ATT	-0.0169	-0.0114	-0.0191
(s.e.)	0.0053	0.0058	0.0058
Pre-trend test (p-value)	0.642	0.507	0.643
Pre-treatment mean productivity	2.806	2.806	2.806
N (firm X year)	75861	75861	75861

Notes: Tabular version of Figure 5. The dependent variable is firm-level TFP estimated following Ackerberg et al. (2015). The table reports the estimated effects of a technological M&A on productivity from: (i) the two-way fixed effects (TWFE) OLS estimator, (ii) Sun and Abraham (2021), and (iii) Callaway and Sant'Anna (2021). Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal; all effects are relative to the year before the deal. The TWFE and Sun and Abraham (2021) specifications include firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, TFP is pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E8: Effect of Technological M&A on Markup, with Controls (Figure D1)

	(1) TWFE OLS	(2) Sun & Abraham	(3) Callaway & Sant'Anna
Year -7	0.0103	0.0012	-0.0043
	(0.0084)	(0.0109)	(0.0084)
Year -6	0.0009	0.0003	-0.0025
	(0.0069)	(0.0077)	(0.0066)
Year -5	-0.0005	-0.0023	-0.0059
	(0.0057)	(0.0064)	(0.0058)
Year -4	-0.0013	-0.0019	0.0003
	(0.0049)	(0.0054)	(0.0049)
Year -3	-0.0016	-0.0012	-0.0010
	(0.0042)	(0.0046)	(0.0042)
Year -2	0.0053	0.0056	0.0031
	(0.0034)	(0.0036)	(0.0034)
Year 0 (deal)	0.0190***	0.0182***	0.0201***
	(0.0040)	(0.0042)	(0.0036)
Year +1	0.0196***	0.0193***	0.0220***
	(0.0046)	(0.0048)	(0.0043)
Year +2	0.0165***	0.0157***	0.0126**
	(0.0056)	(0.0056)	(0.0057)
Year +3	0.0177***	0.0174***	0.0163**
	(0.0059)	(0.0059)	(0.0068)
Year +4	0.0276***	0.0276***	0.0296***
	(0.0069)	(0.0071)	(0.0072)
Year +5	0.0328***	0.0313***	0.0306***
	(0.0077)	(0.0086)	(0.0084)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
Average post-treatment ATT	0.0195	0.0216	0.0215
(s.e.)	0.0046	0.0047	0.0044
Pre-trend test (p-value)	0.145	0.409	0.569
Pre-treatment mean markup	1.383	1.383	1.383
N (firm X year)	71580	71580	71580

Notes: Tabular version of Figure D1. The table reports the estimated effects of a technological M&A on firm markup from: (i) the two-way fixed effects (TWFE) OLS estimator, (ii) Sun and Abraham (2021), and (iii) Callaway and Sant'Anna (2021). Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal; all effects are relative to the year before the deal. Controls include log-transformed firm age and capital stock and IHS-transformed patent stock. The TWFE and Sun and Abraham (2021) specifications include firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A, and the estimates also control for log real operating revenue. Callaway and Sant'Anna (2021) estimates adjust for pre-treatment covariate values using the doubly robust methodology of Sant'Anna and Zhao (2020). Standard errors (displayed in parentheses) are clustered by firm.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E9: Effect of Technological M&A on Markup, Matched Sample (Figure D2)

	(1) TWFE OLS	(2) Sun & Abraham	(3) Callaway & Sant'Anna
Year -7	0.0070	-0.0033	-0.0033
	(0.0102)	(0.0125)	(0.0097)
Year -6	-0.0023	-0.0015	0.0046
	(0.0079)	(0.0097)	(0.0081)
Year -5	-0.0039	-0.0041	-0.0019
	(0.0065)	(0.0080)	(0.0068)
Year -4	-0.0042	-0.0017	0.0027
	(0.0054)	(0.0065)	(0.0056)
Year -3	-0.0013	0.0015	0.0019
	(0.0046)	(0.0054)	(0.0048)
Year -2	0.0047	0.0059	0.0045
	(0.0034)	(0.0040)	(0.0037)
Year 0 (deal)	0.0172***	0.0165***	0.0186***
	(0.0039)	(0.0044)	(0.0040)
Year +1	0.0159***	0.0180***	0.0196***
	(0.0045)	(0.0052)	(0.0047)
Year +2	0.0162***	0.0158**	0.0151**
	(0.0057)	(0.0062)	(0.0060)
Year +3	0.0158**	0.0162**	0.0170**
	(0.0065)	(0.0070)	(0.0069)
Year +4	0.0253***	0.0280***	0.0312***
	(0.0077)	(0.0087)	(0.0087)
Year +5	0.0303***	0.0290***	0.0359***
	(0.0090)	(0.0104)	(0.0101)
Firm FE	Yes	Yes	Yes
Ind X Year FE	Yes	Yes	Yes
Average post-treatment ATT	0.0177	0.0206	0.0228
(s.e.)	0.0047	0.0056	0.0052
Pre-trend test (p-value)	0.186	0.637	0.690
Pre-treatment mean markup	1.383	1.383	1.383
N (firm X year)	24311	24311	24311

Notes: Tabular version of Figure D2. The table reports the estimated effects of a technological M&A on firm markup on the Coarsened Exact Matching (CEM) sample, from: (i) the two-way fixed effects (TWFE) OLS estimator, (ii) Sun and Abraham (2021), and (iii) Callaway and Sant'Anna (2021). Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal; all effects are relative to the year before the deal. The comparison group is matched exactly on the two-digit NACE rev. 2 industry code and on the M&A year (year of first technological M&A for treated firms; year of first M&A of any type for comparison firms), retaining one comparison firm per treated firm. The TWFE and Sun and Abraham (2021) specifications include firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E10: Effect of Technological M&A on Markup, Matched Sample with Controls (Figure D3)

	(1) Callaway & Sant'Anna
Year -7	-0.0058
	(0.0101)
Year -6	0.0012
	(0.0082)
Year -5	-0.0049
	(0.0069)
Year -4	0.0036
	(0.0058)
Year -3	0.0006
	(0.0048)
Year -2	0.0044
	(0.0038)
Year 0 (deal)	0.0174***
	(0.0040)
Year +1	0.0198***
	(0.0047)
Year +2	0.0139**
	(0.0061)
Year +3	0.0170**
	(0.0071)
Year +4	0.0280***
	(0.0090)
Year +5	0.0325***
	(0.0106)
Firm FE	Yes
Ind X Year FE	Yes
Pre-deal covariates	Yes
Average post-treatment ATT	0.0210
(s.e.)	0.0054
Pre-trend test (p-value)	0.419
Pre-treatment mean markup	1.383
N (firm X year)	24311

Notes: Tabular version of Figure D3. The table reports the estimated effects of a technological M&A on firm markup from the Callaway and Sant'Anna (2021) estimator on the Coarsened Exact Matching (CEM) sample, adding pre-treatment covariate controls. Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal; all effects are relative to the year before the deal. The comparison group is matched on the two-digit NACE rev. 2 industry code and the M&A year, retaining one comparison firm per treated firm. Controls are pre-treatment firm-level means of log capital stock, log firm age, and IHS-transformed patent stock computed via the perpetual inventory method with 5 percent depreciation, adjusted for using the doubly robust methodology of Sant'Anna and Zhao (2020). Standard errors (displayed in parentheses) are clustered by firm.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E11: Effect of Technological M&A on Markup, Not-Yet-Treated Controls Only (Figure D4)

	(1) TWFE OLS	(2) Callaway & Sant'Anna
Year -7	-0.0085	-0.0189
	(0.0241)	(0.0123)
Year -6	-0.0139	-0.0069
	(0.0185)	(0.0099)
Year -5	-0.0105	-0.0110
	(0.0150)	(0.0080)
Year -4	-0.0094	-0.0043
	(0.0114)	(0.0066)
Year -3	-0.0060	-0.0011
	(0.0080)	(0.0052)
Year -2	0.0018	0.0048
	(0.0047)	(0.0040)
Year 0 (deal)	0.0201***	0.0184***
	(0.0048)	(0.0045)
Year +1	0.0190**	0.0166***
	(0.0077)	(0.0052)
Year +2	0.0203*	0.0183***
	(0.0112)	(0.0071)
Year +3	0.0206	0.0239**
	(0.0148)	(0.0099)
Year +4	0.0309*	0.0246*
	(0.0184)	(0.0137)
Year +5	0.0372*	0.0286
	(0.0225)	(0.0174)
Firm FE	Yes	Yes
Ind X Year FE	Yes	Yes
Average post-treatment ATT	0.0146	0.0207
(s.e.)	0.0046	0.0061
Pre-trend test (p-value)	0.354	0.256
Pre-treatment mean markup	1.384	1.384
N (firm X year)	13356	13356

Notes: Tabular version of Figure D4. The table reports the estimated effects of a technological M&A on firm markup when the comparison group comprises only not-yet-treated firms, from: (i) the two-way fixed effects (TWFE) OLS estimator and (ii) Callaway and Sant'Anna (2021). Coefficients are shown for up to seven years before the deal, the deal year, and up to five years after the deal; all effects are relative to the year before the deal. The TWFE specification includes firm and industry-by-year fixed effects; for the Callaway and Sant'Anna (2021) estimator, markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table E12: Effect of Technological M&A on Markup, Robust to Reversals (Figure D5)

	(1) dCdH (2024)
Year -6	0.0164
	(0.0101)
Year -5	0.0016
	(0.0069)
Year -4	0.0016
	(0.0055)
Year -3	0.0007
	(0.0042)
Year -2	0.0026
	(0.0033)
Year 0 (deal year)	0.0208***
	(0.0036)
Year +1	0.0212***
	(0.0042)
Year +2	0.0167***
	(0.0052)
Year +3	0.0196***
	(0.0057)
Year +4	0.0309***
	(0.0070)
Year +5	0.0319***
	(0.0080)
Firm FE	Yes
Year FE	Yes
Average post-treatment ATT	0.0235
(s.e.)	0.0042
Placebo test (p-value)	0.597
Pre-treatment mean markup	1.386
N (firm X year)	53990

Notes: Tabular version of Figure D5. The table reports the estimated effects of a technological M&A on firm markup under the non-absorbing treatment definition, using the heterogeneity-robust event-study estimator of De Chaisemartin and d'Haultfoeuille (2024), which compares switchers—firms whose treatment status changes during the panel—to firms sharing the same baseline treatment that have not yet switched. All effects are relative to the year before the deal. The reported placebo test is the joint Wald test of the null that all five placebo estimates are zero. Markups are pre-residualized on industry-by-year means of firms with no technological M&A. Standard errors (displayed in parentheses) are clustered by firm.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The UNU-MERIT WORKING Paper Series

2026-01 *Beyond Monetary Measures: Determinants of Multidimensional Poverty among Vulnerable Older Children and Youth in East Germany. A Perspective from Capability Approach and Capital Theory* by Charlotte Margarete Jost & Zina Nimeh
2026-02 *Sustainable urban rail integration: the case of the city of Maastricht* by Jo Coenen and Luc Soete
2026-03 *Redirect investment to stimulate the economy* by Huub Meijers and Joan Muysken
2026-04 *Airless democracy: Air pollution and voter turnout* by Giulia Rossello, Maria Antonietta Reatini, Gabriele Pinto, and Giorgio Cattani
2026-05 *Beyond size: Unpacking the effects of small-scale de-reservation policy on innovation and performance of manufacturing plants in India* by George Paily, Nanditha Mathew and Carlo Pietrobelli
2026-06 *Deconstructing Quality Infrastructure: A Quantitative Analysis of the Global QI Index* by Ulrich Harmes-Liedtke, Carlo Pietrobelli, Ann-Sara Ramkissoon and Francesco Salustri
2026-07 *The cost of bureaucratic fragmentation: Business tax evasion and revenue mobilization in a low-income country* by Stephan Dietrich, Yannick Markhof and Rose Camille Vincent
2026-08 *Gender and Shock Responsiveness in Social Protection Systems: Lessons from Jordan* by Zina Nimeh, Chris de Neubourg, Míriam Carrera and Nesha Ramful
2026-09 *Automation and Human Capital Investment* by Sidharth Rony
2026-10 *LLM Meets Job Advertisements: Unmasking Skill Premia in the UK* by Sidharth Rony
2026-11 *Gendering green transition in Africa: Implications of transformative policies* by Gideon Ndubuisi and Elvis Korku Avenyo
2026-12 *The interplay between AI and technological relatedness in shaping regional innovation in Europe* by Francesco D'Alessandro, Enrico Santarelli and Marco Vivarelli
2026-13 *Patents, Pay and People: How Innovation Reshapes the* Workforce by Çiğdem Ekiz, Eren Gürer and Erol Taymaz
2026-14 *Age, Education, and Occupational Gradients in Artificial Intelligence Anxiety* by Klaus F. Zimmermann, Gökhan Karabulut, Zeynep Keklik and Kıvanç Altıntaş
2026-15 *Partisan Demographics* by Damilola Afolabi and Klaus F. Zimmermann
2026-16 *Global Value Chains in the Era of Industry 4.0 and AI: Extractive Industries in Latin America* by Carlo Pietrobelli and Constanza Araya-Ibarra
2026-17 *Superstars Have the Power? Technological M&As and Firm Markups* by Cecilia Seri1, Elena Zaurino, Maria Martinez Cillero, Lorenzo Napolitano and Francesco Rentocchini