

Digital (Killer?) Acquisitions*

Florian Ederer[†] Regina Seibel[‡] Timothy Simcoe[§]

September 11, 2025

Abstract

This paper examines the effects of 1,200 acquisitions by major technology firms on innovation. Using detailed patent and workforce data linked to technology acquisitions and a suite of event-study and difference-in-differences designs, we document four main findings. First, although most acquired startups hold no patents, those with patents tend to operate in technology areas where the acquirer already has a presence and which subsequently experience further acquisition activity. Second, innovation typically rises before an acquisition but only persists afterward when further acquisitions occur, suggesting that acquisitions respond to innovation trends rather than initiating them. Third, at the patent level, acquired patents receive significantly more citations after the acquisition than comparable patents, and these effects are not driven solely by self-citations from the acquiring firm. These post-acquisition citation effects are smaller when more employees from the acquired firm are retained, consistent with innovation spillovers occurring through employee mobility. Fourth, we document significant workforce attrition exceeding 50% on average at target firms three years post-acquisition. Our results suggest that acquisitions by digital incumbents often amplify, rather than suppress, the diffusion and visibility of acquired technologies.

JEL Classification: L41, O31, O34

Keywords: Innovation, Acquisitions, Patents

*We thank Heski Bar-Isaac, EL Hadi Caoui, Joshua Gans, Alexander Hoagland, Mitsuru Igami, Hong Luo, Matt Mitchell, Mario Leccese, Zi Yang Kang and seminar audiences at University of Toronto for helpful comments. Ariel Hsieh, Rachelle Lang, Nathan Shekita, Meg Snyder, and Davin Wang provided outstanding research assistance.

[†]Boston University Questrom School of Business, CEPR, ECGI, and NBER, ederer@bu.edu

[‡]Toronto Rotman School of Management, regina.seibel@rotman.utoronto.ca

[§]Boston University Questrom School of Business and NBER, tsimcoe@bu.edu

1 Introduction

The rising dominance of large digital platforms has sparked renewed interest in the competitive implications of their acquisition behavior. Over the past two decades, leading technology firms—including Amazon, Apple, Facebook, Google, and Microsoft (commonly referred to as GAFAM)—as well as Cisco, Intel, and Qualcomm have each engaged in hundreds of acquisitions, often targeting early-stage startups in adjacent or emerging technology markets. These transactions raise important questions about how markets for technology function, and whether acquisitions facilitate or hinder the commercialization and diffusion of innovative ideas.

Given the complementary assets required to bring new technologies to market, innovative startups often rely on licensing deals or acquisitions to appropriate returns from their inventions (Teece, 1986; Gans and Stern, 2000; Gans, Hsu and Stern, 2002; Arora, Fosfuri and Gambardella, 2004). At the same time, regulators, policymakers, and scholars have raised concerns that some of these acquisitions may serve anticompetitive purposes. A growing body of research emphasizes innovation-based theories of harm, highlighting that mergers can reduce innovation incentives, eliminate nascent competition, or alter the direction of technological progress (Federico, Langus and Valletti, 2017, 2018; Federico, Scott Morton and Shapiro, 2020; Cunningham, Ederer and Ma, 2021; Valletti, 2025). These concerns are particularly salient in digital markets, where acquisitions can affect not only current competition but also the trajectory of future innovation. Recent policy discussions and interventions, including the 2023 U.S. Merger Guidelines and the EU Digital Markets Act, explicitly recognize that repeated or serial acquisitions may cumulatively undermine competitive innovation dynamics.¹

This paper investigates the effects of startup acquisitions by major digital incumbents on subsequent innovation outcomes within narrowly defined technological domains. Instead of focusing on the ex-ante incentives of potential acquisition targets – a common emphasis in the theoretical literature² – we examine ex-post outcomes using an event-study framework. Specifically, we analyze changes in patenting and patent citations before and after an acquisition, thereby assessing how these deals reshape the trajectory of innovation in the affected

¹Concerns about acquisitions that may lead to the termination of overlapping innovation projects were first raised in the pharmaceutical industry (Cunningham et al., 2021), but have since been applied to the digital economy (OECD, 2020; Federal Trade Commission, 2021) and other R&D-based industries. However, in contrast to pharmaceuticals the innovation environment in the tech industry is less reliant on formal intellectual property and more centered around talent acquisition and platform integration (Gawer, Cusumano et al., 2002; Jin, Leccese and Wagman, 2023; Olmsted-Rumsey, Puglisi and Wu, 2024).

²See, in addition to the previously mentioned contributions, Kamepalli, Rajan and Zingales (2020), Fumagalli, Motta and Tarantino (2022), Bryan and Hovenkamp (2020), Motta and Shelegia (2025), and Letina, Schmutzler and Seibel (2024).

technology space.

We make three main contributions to the literature on acquisitions, innovation, and competition policy. First, we provide a comprehensive empirical mapping of the acquisition behavior of eight major digital firms (the GAFAM firms plus Cisco, Intel and Qualcomm). We manually compile a dataset of 1,200 acquisitions and link this information to patent data from PatentsView and labor market data from Revelio Labs. We show that while these firms have acquired a large number of startups, the majority of targets do not hold patents at the time of acquisition which suggests that many deals are motivated by access to capabilities, products, customers, or personnel rather than formal intellectual property. We also find that while retention rates vary substantially across deals, on average, only 45% of target firm employees remain at the acquirer three years after the acquisition, highlighting the fragility of workforce integration.

We find that acquisitions are highly concentrated in specific technological domains and often represent repeated engagement by the same acquirer in a given area. For instance, firms like Intel and Qualcomm which are traditionally hardware-focused, tend to acquire targets in Cooperative Patent Classification (CPC) categories where they already hold patents, suggesting a form of cumulative innovation or vertical integration. In contrast, software-oriented firms such as Facebook or Amazon are more likely to pursue acquisitions without patents, perhaps reflecting a focus on user bases, engineering talent, or data assets.

Second, we construct a panel dataset at the technology-class-year level that allows us to analyze innovation trends surrounding each deal. We find that innovation frequently precedes acquisition, a finding that undermines the “killer acquisition” narrative in its strongest form, where the incumbent preemptively shuts down nascent rivals. Instead, we find that many acquisitions follow a surge of patenting within the same CPC groups where the acquired startup obtains patents, consistent with incumbents acquiring startups to access promising technologies.

Using a suite of event-study and difference-in-differences analyses, we also estimate the impact of acquisitions on subsequent innovation activity. We find that patenting activity in affected CPC groups generally continues to rise after an acquisition. This pattern is robust to matching treated CPC groups with control groups based on pre-acquisition growth trends and controlling for a range of fixed effects. However, the post-acquisition growth in patenting is only present when follow-on acquisitions occur in the same technology space. Recent antitrust policy has also shifted toward recognizing the broader strategic implications of repeated acquisitions. Rather than evaluating transactions solely on a case-by-case basis, regulators are increasingly concerned with the cumulative effects of serial acquisitions within a particular technological or geographic domain. The 2023 U.S. Merger Guidelines explicitly

acknowledge that mergers occurring as part of a pattern or roll-up strategy may be assessed collectively rather than in isolation (U.S. Department of Justice and Federal Trade Commission, 2023). This perspective is reinforced by recent empirical work on acquisition rollups in health care markets, which documents the anticompetitive risks posed by fragmented yet systematic buying strategies (Asil, Ramos, Starc and Wollmann, 2024). Our analysis of repeated acquisitions within CPC technology classes provides complementary evidence in the digital sector and speaks directly to this evolving regulatory framework.

Third, in addition to CPC group-level analyses, we examine the effects of acquisitions at the level of individual patents. Using a difference-in-differences framework, we show that acquired patents receive significantly more citations after the acquisition than comparable patents in the same CPC group. These effects are not limited to citations from the acquiring firm, but extend to third-party citations, suggesting that acquisitions increase the visibility and influence of acquired technologies across the broader technology ecosystem. The post-acquisition citation boost is especially pronounced for first-acquired targets in CPC groups that experience follow-on acquisitions, indicating that the strategic sequencing of acquisitions may amplify knowledge diffusion. We also find that post-acquisition citation effects are significantly smaller when a larger share of the workforce is retained. This surprising result suggests that innovation spillovers may be driven in part by departing employees, who disseminate knowledge beyond the acquiring firm.

Our findings complicate the popular narrative that digital acquisitions by large incumbents are predominantly anti-competitive. While we do not rule out the possibility of killer acquisitions in individual cases, our evidence suggests that most deals by GAFAM and related firms are motivated by complementarity rather than suppression. Innovation tends to rise rather than fall in the wake of acquisition, particularly in technological domains where the acquirer has prior experience and continues to make follow-on investments.

This paper contributes to several strands of literature at the intersection of innovation, acquisitions, and competition policy. First, we build on a growing body of empirical work on killer acquisitions, beginning with Cunningham et al. (2021) in the pharmaceutical industry. Their framework identifies acquisitions that lead to the termination of overlapping drug development projects, thereby suppressing future competition. Subsequent research has extended this idea to other sectors. For example, Lemos and Resende (2023) examine digital markets, while Kamepalli et al. (2020) explore data-driven barriers to entry following acquisitions by dominant platforms. Our analysis complements this work by focusing on ex-post outcomes in patenting activity rather than ex-ante project terminations. Jin et al. (2023) show that GAFAM made more tech acquisitions per firm than other top acquirers, targeting younger and more consumer-facing firms. Crucially, they find no evidence that

GAFAM acquisitions reduce entry by other acquirers in the same category, suggesting that such deals do not deter competition, thereby complementing our finding that, on average, acquisitions do not suppress innovation within targeted technologies.

Second, we relate to a broader literature on the link between competition policy and innovation (Gilbert, 2020). Buehler and Schmutzler (2005), Valletti and Zenger (2017), and Gaffard and Quatraro (2022) provide theoretical and empirical perspectives on how consolidation can affect both the intensity and direction of innovative effort. Our event-study design allows us to trace such effects dynamically within narrowly defined technological areas. In contrast to most existing work, we focus specifically on startup acquisitions by digital incumbents and analyze repeated treatments over time. Like us, Watzinger, Fackler, Nagler and Schnitzer (2020) and Poege (2024) use patent data to study how competition policy influences innovation across a set of related technology domains. Whereas they focus on a single historically significant enforcement action (the 1956 AT&T consent decree and the post-war breakup of IG Farben, respectively), we consider a large number of more recent acquisitions.

Third, we contribute to a broad literature on markets for technology (e.g., Arora and Gambardella, 2010, *inter alia*), that frequently uses patent data to measure innovation outcomes. The use of Cooperative Patent Classification (CPC) groups as a proxy for technological domains follows prior work such as Hall, Jaffe and Trajtenberg (2001) and Bloom, Schankerman and Van Reenen (2013). We also draw on insights from the literature on patent citations as a measure of knowledge spillovers, including Jaffe, Trajtenberg and Henderson (1993) and Akcigit, Grigsby and Nicholas (2016).

The rest of the paper is organized as follows. Section 2 describes our data sources and the construction of the CPC-year panel. Section 2 provides descriptive statistics on acquisitions, targets, and patenting patterns. Section 4 presents our main empirical analysis of post-acquisition patenting outcomes and uses citation-based metrics to capture the influence and quality of innovation. Section 5 concludes with a discussion of implications for competition policy and future research directions.

2 Data

To investigate the effects of acquisitions by major digital firms on innovation, we construct a novel dataset that combines acquisition activity with detailed patent-level information. This section describes our data sources, the construction of key variables, and the process by which we merge acquisition records with patent data.

2.1 Acquisition Data

We begin by assembling a manually curated dataset of startup acquisitions by eight prominent technology companies: Amazon, Apple, Cisco, Facebook, Google, Intel, Microsoft, and Qualcomm. We selected these firms due to their longstanding presence in the digital and information technology sectors, their high frequency of acquisition activity, and their central role in current debates on digital competition and innovation policy. While Amazon, Apple, Facebook, Google, and Microsoft are typically characterized as digital platforms serving end users directly, firms like Intel, Cisco, and Qualcomm are equally important players, operating as key infrastructure and input providers in the broader digital ecosystem. These firms specialize in semiconductors, networking hardware, and communication protocols, and their acquisitions often target upstream innovations essential for enabling broader technological progress.

All eight firms have drawn antitrust scrutiny for their acquisition behavior. Notable examples include the FTC’s investigation of Facebook’s acquisitions of Instagram and WhatsApp,³ the U.S. Department of Justice’s (DOJ) antitrust lawsuit against Apple alleging monopolistic practices in the smartphone market,⁴ and the FTC’s inquiries into Microsoft’s, Google’s and Amazon’s investments in generative AI companies like OpenAI and Anthropic.⁵ Qualcomm’s proposed acquisition of NXP Semiconductors faced global regulatory delays before collapsing in 2018,⁶ while Intel and Cisco have also faced merger reviews due to their active roles in consolidating upstream technology markets. In addition, the European Union’s Digital Markets Act (DMA), obliges designated “gatekeepers” (including all five GAFAM firms) to notify the European Commission of any planned acquisition involving a digital sector company, even if the deal does not meet the formal merger notification thresholds at the EU or national level.⁷ Our data capture this diverse spectrum of acquisition strategies and regulatory attention across platform and infrastructure firms alike.

³Federal Trade Commission, “FTC Sues Facebook for Illegal Monopolization,” December 9, 2020, <https://www.ftc.gov/news-events/news/press-releases/2020/12/ftc-sues-facebook-illegal-monopolization>.

⁴U.S. Department of Justice, “Justice Department Sues Apple for Monopolizing Smartphone Markets,” March 21, 2024, <https://www.justice.gov/archives/opa/gallery/justice-department-sues-apple-monopolizing-smartphone-markets>.

⁵Federal Trade Commission, “FTC Launches Inquiry into Generative AI Investments and Partnerships,” January 25, 2024, <https://www.ftc.gov/news-events/news/press-releases/2024/01/ftc-launches-inquiry-generative-ai-investments-partnerships>.

⁶Reuters, “Qualcomm Ends \$44 Billion NXP Bid After Failing to Win China Approval,” July 26, 2018, <https://www.reuters.com/article/technology/qualcomm-ends-44-billion-nxp-bid-after-failing-to-win-china-approval-idUSKBN1KF18X/>.

⁷Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector (Digital Markets Act), <https://eur-lex.europa.eu/eli/reg/2022/1925/oj>.

For each acquisition, we record the name of the acquired company, the year in which the acquisition occurred, and, where available, information about the target’s patenting activity and deal value. In order to focus on startup acquisitions, which we define as acquisitions involving relatively small and early-stage companies, we exclude large deals involving targets with extensive patent portfolios (e.g., Apple’s acquisition of Intel’s modem business). These exclusions allow us to avoid confounding effects from horizontal consolidation and instead focus on the acquisition of emerging innovators. We identify 1,190 acquisitions of such emerging innovators, comprising the vast majority of all acquisitions by the eight technology companies under study.

We link each acquired startup to its portfolio of patents using name-based matches between targets and assignees in the U.S. patent data. This matching process enables us to identify the specific technological areas within the Cooperative Patent Classification (CPC) system where the acquired firm was active prior to the acquisition.

2.2 Patent Data

To characterize technological activity, we rely on data from the PatentsView database (USPTO 2016). Specifically, we extract data on patent filings and grants, including application years and Cooperative Patent Classification (CPC) codes. We also obtain information on patent assignees, which allows us to link patents to the firm that holds the right at the time of issuance.

Throughout the paper, we use CPC codes to group together patents in related technology areas. The Cooperative Patent Classification system is a hierarchical system devised by the US Patent and Trademark Office (USPTO) and the European Patent Office (EPO). The main purpose of the CPC system is to facilitate prior art searches by patent examiners, so the patents within a code correspond to similar technologies (though we would not go so far as to argue that they constitute technology markets in the sense of having a similar level of substitutability among patents). All US patents are assigned at least one primary CPC code, and possibly several secondary codes. The system has over 250,000 distinct categories, which allows for billions of possible combinations. Throughout this paper, we refer to the six-digit main group associated with a patent’s primary CPC code as that patent’s “CPC class.”

To focus our analysis on technology areas most relevant to the digital sector, we restrict attention to patents classified under CPC sections G (Physics) and H (Electricity) with 1,289 CPC classes. These two sections encompass the majority of innovation activity in

information technology, software, electronics, and telecommunications.⁸

2.3 Data on Employee Work Histories

We use data from Revelio Labs to analyze employee-level career trajectories before and after acquisition events. Revelio Labs aggregates publicly available resumes, primarily from LinkedIn, to construct standardized longitudinal records of individuals' employment histories. This dataset allows us to trace job transitions, employer affiliations, and role durations at the individual level, making it well-suited to study labor market dynamics around acquisitions.

To link employees to acquired firms, we begin with our list of 1,190 startup acquisitions and search for exact matches between target firm names and employer names in the Revelio data.⁹ Because Revelio data is sourced largely from LinkedIn, which was founded in 2003, and because we require pre-acquisition employment histories to evaluate changes over time, we limit the sample to acquisitions occurring in 2005 or later. This restriction yields 932 acquisitions. Of these, we manage to cleanly identify and verify 606 targets in the Revelio data. We similarly identify the acquiring firms in the Revelio data. In doing so, we map not only to the parent-acquiring entity but also include all affiliated subsidiaries. For instance, positions listed at Instagram, DeepMind, or GitHub are treated as positions at Facebook, Google, and Microsoft, respectively, after their respective acquisition dates.

The resulting data provides a dynamic view of workforce changes surrounding acquisition events, allowing us to estimate retention patterns, employee turnover, and the relationship between post-acquisition integration and innovation outcomes. This labor market lens complements our patent-based analyses and offers a more complete picture of the organizational consequences of digital acquisitions.

2.4 Dataset Construction

Our empirical analysis uses two panel datasets. In the first dataset, each observation represents a unique technological field (CPC class) in a specific year. For each CPC-year combination, we calculate the number of patent applications filed, identify whether an acquisition occurred involving a target firm that had previously patented in that CPC group, and record

⁸While this restriction excludes a small number of patents filed by acquirers or targets in our sample, it screens out a large number of CPC classes that are never used by firms in our sample, and it dramatically reduces computational burdens. Across the entire corpus of US patents, there are 7,005 CPC classes, and each class has an average 30.3 new patent applications per year.

⁹Revelio Labs also offers matched patent data, which we use to check and enhance the matching of patent-holding targets.

which of the eight focal technology firms carried out the acquisition. This structure allows us to examine changes in innovation activity (measured by patenting intensity) before and after acquisition events within the same technological area. In addition, it enables us to compare treated CPC groups to those that were never subject to an acquisition, thereby facilitating a difference-in-differences and event-study identification strategy.

We construct several variables to characterize the acquisition behavior of each firm and the structure of innovation within each CPC group. These include indicators for whether the acquirer had patented in the CPC group prior to the acquisition, whether the acquisition represented the first such transaction in that technology area, and whether multiple acquisitions occurred in the same CPC group over time. These variables allow us to explore patterns of repeated acquisitions, the role of technological expertise, and the persistence of innovation effects following a merger.

In the second panel, each observation represents a single patent in a given year. We restrict the sample to patents in CPC groups containing at least one patent filed by a target firm before its acquisition, and track whether each patent was filed by the acquired firm or a third-party. Our main outcome variable is the number of prior-art citations that a focal patent receives from other U.S. patents filed or granted in a particular year. These forward citations are a widely used measure of the economic and technological significance of the cited patent, and the panel structure allow us to trace the relative importance of specific patents through time, before and after each acquisition event.

These two datasets provide a foundation for analyzing how acquisitions influence innovation in targeted technology areas. By combining patent data with manually collected acquisition records, we are able to track the timing, intensity, and nature of innovation — at both the broad technology and the individual patent level — both before and after each acquisition event. This structure enables us to evaluate the relationship between acquisition activity and technological development with a relatively high degree of granularity.

3 Descriptive Statistics

3.1 Big Tech Acquisitions

Table 1 provides an overview of the acquisition activity by each of the eight major digital firms included in our analysis, along with summary statistics on the patent holdings of their targets. The first column shows the number of acquisitions for each acquirer in our sample. The second shows that share of targets with at least one patent. For example, just 23% of Amazon’s targets had applied for a patent prior to acquisition, but those patenting targets

owned 10 patents in four different CPC groups. This means that for Amazon, on average, one acquisition of a patent-holding target leads to 4 CPC-level “treatments” in the event study models described below. Comparing the various acquirers several patterns emerge. First, across all acquirers, a substantial share of startup targets do not hold any patents at the time of acquisition. The probability that an acquired firm holds at least one patent ranges from as low as 17% for Facebook to as high as 70% for Qualcomm. This heterogeneity likely reflects differences in acquisition strategies, with some firms (particularly those in semiconductors and hardware-intensive sectors) placing greater emphasis on acquiring patented technologies.

Second, conditional on acquiring a target with patents, the average number of patents also varies across firms. Intel and Qualcomm stand out with average target portfolios exceeding 20 patents per acquisition. These targets also span a wider technological scope, as indicated by the number of distinct CPC codes associated with their patents. In contrast, targets acquired by Facebook and Amazon typically have smaller patent portfolios and fewer CPC categories, suggesting a focus on capabilities beyond formal intellectual property such as talent, products, or user data.

The next column of Table 1 reports the average number of G- and H-class CPC codes per patent-holding acquisition. These CPC sections (Physics and Electricity, respectively) account for the bulk of digital and information technology innovation. The similarity between the average total CPC codes and G&H-specific CPC codes suggests that most patent-relevant acquisitions in our sample are concentrated within these core technological domains. This suggests that restricting the analysis to G&H CPC codes represents a minor but targeted refinement of the full sample.

The last three columns in Table 1 provide descriptive statistics on post-acquisition retention of employees from the acquired firm. On average, around 50 percent of employees have left the acquirer by three years after the merger, and roughly two-thirds have departed within 5 years. The Retention Gap statistic in the final column is the difference in retention rate between employees of the acquired firm, and the retention rates for a matched sample of other employees who joined the acquiring firm in the same year, over a three-year period. Interestingly, the Retention Gap varies across acquirers by more than the retention rate, which suggests substantial differences in the underlying job separation rate at these large digital firms.

Overall, the data in Table 1 highlight the diversity of acquisition patterns across large digital firms and underscore the importance of distinguishing between acquisitions of patent-rich and patent-poor targets. These patterns also motivate our event-study approach, which tracks innovation outcomes in narrowly defined technological fields before and after acquisition events.

Table 1: Number of targets, patents acquired, and retention rates by each company

Company	#Acquisitions	Pr(AnyPatent)	#Patents	#CPC	#CPC (G&H)	RR 3y	RR 5y	RGap 3y
Amazon	103	0.23	10.38	4.17	3.04	0.50	0.39	0.23
Apple	126	0.36	16.44	4.93	4.22	0.56	0.49	0.19
Cisco	222	0.35	12.17	4.12	3.91	0.50	0.36	0.05
Facebook	102	0.17	3.71	2.29	2.12	0.46	0.33	0.44
Google	260	0.24	13.24	3.16	2.76	0.45	0.36	0.38
Intel	97	0.57	21.40	7.38	7.07	0.47	0.38	0.16
Microsoft	236	0.31	14.74	4.12	3.95	0.44	0.33	0.23
Qualcomm	44	0.70	20.74	7.61	7.35	0.54	0.40	0.27

NOTES: The means for #Patents and #CPC are conditional on target firm holding at least one patent. RR 3y RR 5y denotes denote, respectively, the 3 and 5-year post-acquisition employee retention rate. Retention Gap equals the difference in post-acquisition retention rate between acquiring and target firm for employees hired in same year.

Table 2 reports median deal values (in millions of USD) for acquisitions by each of the eight focal technology firms, separately for targets with and without patent holdings.¹⁰ The results reveal a clear pattern: across nearly all firms, acquisitions of patent-holding targets are associated with substantially larger transaction values. These findings reinforce the importance of distinguishing between acquisitions that involve patented innovation and those that do not. They also suggest that while many acquisitions are motivated by access to capabilities or personnel, patents remain a key driver of valuation and strategic interest for big tech acquirers.

Table 2: Median deal values (in millions of USD)

Company	Targets without patents	Targets with patents
Amazon	86	300
Apple	30	239
Cisco	126	282
Facebook	40	300
Google	61	500
Intel	163	250
Microsoft	128	200
Qualcomm	26	190
Total	100	259


For instance, Apple paid roughly eight times more for the median patent-holding target (\$239 million) than for the median target without patents (\$30 million). Similarly, Google’s median deal size increases from \$61 million to \$500 million when the target holds patents. This suggests that patent holdings may serve as a signal of technological quality, maturity, or strategic value, especially in software-oriented acquisitions. Even for firms like Intel and Qualcomm which tend to operate in patent-intensive domains, deal values are markedly higher when patents are present.¹¹

Figure 1 illustrates the annual frequency of startup acquisitions over the sample period. The figure serves two key purposes. First, it documents the overall volume and dynamics of acquisition activity across years. Second, it provides visual context for understanding the timing of treatment events in our event-study framework.

¹⁰We focus on median rather than mean deal value because the latter statistics is sensitive to outliers.

¹¹This patent premium is similarly large when taking the size of the target into account. In Table A.4 in Appendix A.3, we show that deal values are much larger when patents are present on average, even after controlling for the number of target firm employees. In fact, having more employees seems to amplify the premium paid for targets with patents.

Figure 1: Acquisition Time Trends




The data show that acquisition activity is unevenly distributed over time, with noticeable peaks in certain years. These surges often coincide with broader trends in the technology sector, such as funding booms, IPO droughts, or platform expansion phases. For example, there is a pronounced increase in acquisition activity during the early 2010s, which corresponds to a period of aggressive growth by firms like Google, Facebook, and Amazon. The observed variation in acquisition timing underscores the need to control for year fixed effects in our empirical specifications. It also highlights the importance of matching treated and untreated technological areas based on pre-trends in innovation, as different cohorts of acquisitions may have been subject to distinct macroeconomic or regulatory conditions.

3.2 Acquired Patent CPC Codes

From this point forward, we restrict our analysis to 1,289 CPC codes within sections G (Physics) and H (Electricity). This restriction is analytically useful because it concentrates our attention on core technology areas most relevant to the digital economy, including elec-

Figure 2: Number of acquisitions in treated CPC codes



tronics, telecommunications, and software.¹² In total, 303 distinct G&H CPC codes are affected by at least one acquisition in our sample. Among these treated CPC groups, approximately 38% experience only a single acquisition event, while the remaining 62% are treated multiple times, reflecting repeated acquisition activity within the same technological field. This recurrence suggests that certain technology areas serve as ongoing targets of strategic interest for digital incumbents.

Figure 2 provides a histogram of the number of acquisition events in each CPC code, conditional on the group being “treated” at least once. The distribution is highly skewed: while many CPC codes experience only one or two acquisition events, a non-negligible number of technology areas see repeated acquisition activity, with some codes linked to as many as ten acquired firms over the sample period. This pattern helps motivate our investigation of repeated treatments and acquisition waves.

Table 3 examines the technological experience of each acquirer at the time of their first acquisition within a given CPC code. For each firm, we report the number of such first-time acquisition events, the number of cases in which the acquirer had previously patented in the same CPC group, and the share of total patents held by the acquirer within those groups.

The table shows that acquirers frequently enter a CPC group where they have a foun-

¹²As shown earlier in Table 1, this focus does not meaningfully reduce the breadth of our sample. The average number of G&H CPC codes affected per acquisition is very similar to the average across all CPC codes, indicating that most patent-relevant startup acquisitions already fall within these sections.

Table 3: Previous patent experience for Acquirer’s first acquisition in CPC

Company	# Events	# Prev. patent	Patent share	Min	Max
Amazon	50	44	0.0060	0.0001	0.0264
Apple	107	99	0.0096	0.0001	0.0855
Cisco	100	82	0.0097	0.0000	0.0762
Facebook	23	20	0.0040	0.0002	0.0292
Google	97	89	0.0070	0.0002	0.0447
Intel	177	167	0.0223	0.0006	0.1545
Microsoft	103	98	0.0260	0.0003	0.1482
Qualcomm	133	123	0.0170	0.0001	0.1200

dation of prior patenting activity. For example, Intel already owned patents in 167 out of 177 CPC classes that they “entered” through acquisition, and Microsoft had prior patents in 98 of 103 cases. Even for acquirers more typically associated with software and platform services (e.g., Google and Apple) there is a strong tendency to acquire in domains where they already hold intellectual property.

The final three columns of Table 3 report the average share of patents in each CPC group that are held by the acquirer prior to the acquisition, along with the minimum and maximum values of this share across CPC groups. These shares are generally small in absolute terms (often below 3%), but they vary meaningfully across firms. Microsoft and Intel exhibit particularly high average shares, suggesting more entrenched innovation activity within the technological domains they target. At the other end of the spectrum, Amazon and Facebook tend to have lower pre-acquisition shares, consistent with a more exploratory or externally driven innovation strategy.

3.3 Repeat Acquisitions

Given that many CPC groups experience multiple acquisition events, it is informative to examine the structure and sequence of these acquisition waves. Table 4 reports transition probabilities for acquisition events. The table provides a descriptive foundation for understanding how acquisition events cluster within technological domains and motivates our empirical focus on acquisition waves and their implications for post-merger innovation.

Specifically, for each acquirer, we report the probability that an acquisition event is either not followed by any further acquisition in the same CPC group (column “None”) or is followed by an acquisition by any of the eight focal acquirers. Note that the unit of

Table 4: Conditional probabilities of follow-on acquisitions

	#Events	None	Amazon	Apple	Cisco	Facebook	Google	Intel	Microsoft	Qualcomm
Amazon	31	0.48	0.10	0.19	0.10	0.00	0.06	0.26	0.19	0.06
Apple	90	0.32	0.07	0.19	0.12	0.03	0.11	0.26	0.09	0.08
Cisco	91	0.19	0.05	0.05	0.36	0.00	0.11	0.40	0.30	0.32
Facebook	13	0.08	0.08	0.15	0.38	0.15	0.15	0.00	0.38	0.31
Google	85	0.58	0.08	0.08	0.07	0.01	0.11	0.12	0.15	0.08
Intel	168	0.28	0.05	0.20	0.19	0.04	0.12	0.25	0.14	0.23
Microsoft	91	0.21	0.09	0.16	0.22	0.08	0.19	0.32	0.24	0.10
Qualcomm	107	0.27	0.03	0.10	0.24	0.00	0.08	0.31	0.16	0.23
Multiple	83	0.16	0.18	0.23	0.37	0.10	0.22	0.51	0.30	0.24

observation is an acquisition-CPC group pair, so a single acquisition spanning multiple CPC groups will contribute multiple events.

The row labeled “Multiple” summarizes transition probabilities in cases where multiple acquirers make acquisitions in the same CPC group in the same year. These simultaneous events are relatively common and may indicate areas of heightened strategic interest or overlapping innovation trajectories. In these cases, Intel and Microsoft are the most likely firms to appear as follow-on acquirers, again pointing to their broad and sustained involvement across many technology domains.

Several additional insights emerge. First, firms differ markedly in the likelihood that an acquisition event is followed by another in the same CPC group. For example, 58% of Google’s events and 48% of Amazon’s are not followed by any subsequent acquisition, while Cisco and Microsoft are much more likely to engage in repeated activity. Second, follow-on acquisitions often come from the same firm, but not always. Some firms (e.g., Intel) frequently acquire in CPC groups where other firms have already made acquisitions, suggesting either shared interests in certain technologies or strategic responses to rivals’ moves.

Figure 3 presents a histogram of the number of years between successive acquisition events within the same CPC group, conditional on a follow-on acquisition occurring. The figure captures the temporal clustering of acquisitions and helps to characterize the dynamics of repeated activity within a technological domain. The distribution is highly skewed toward short intervals between events. More than 40% of acquisitions happen in the same calendar year as another acquisition in the same CPC group, typically by a different firm. Additionally, over 20% of acquisition events are followed by another within just one year. These short intervals suggest that certain CPC groups experience concentrated bursts of acquisition activity, which may reflect heightened competitive interest or simultaneous recognition of

Figure 3: Years to next acquisition within same CPC group

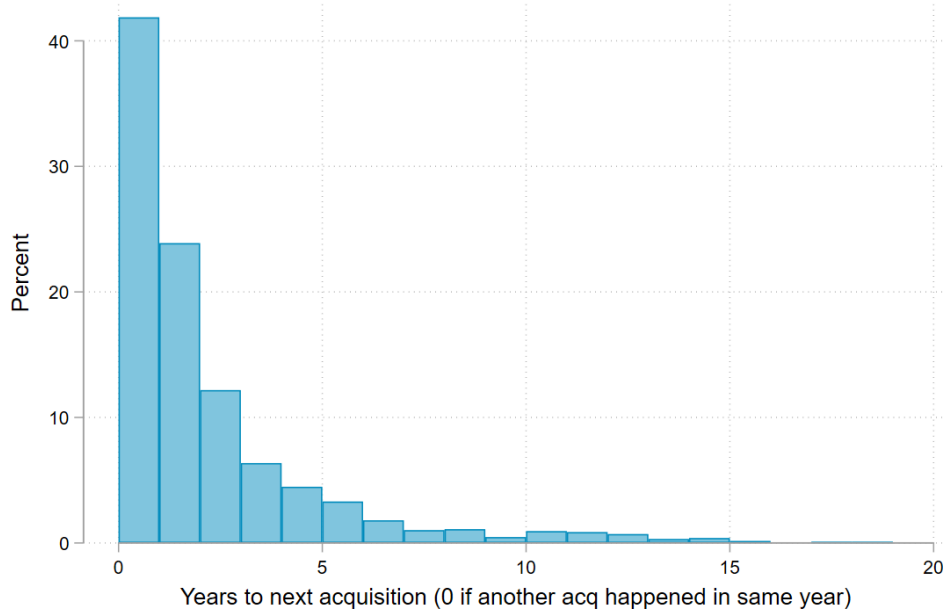


Table 5: Follow-on Acquisition and Wave Positions by Experience

Experience	# Events	No follow-on acq	First in acq wave	Last in acq wave	Mid in acq wave
No	69	0.43	0.27	0.13	0.18
Yes	722	0.10	0.19	0.14	0.56

strategic value across multiple firms.¹³

Table 5 examines how an acquirer’s prior patenting experience in a CPC group relates to the likelihood of follow-on acquisitions in that same technological domain. The results reveal a striking pattern: when the acquirer had already patented in the CPC group before the acquisition, 90% of first acquisitions are followed by additional acquisitions in that group. By contrast, when the acquirer had no prior experience in the CPC group, the first acquisition is much more likely to be an isolated event. This suggests that acquirers with established technological footholds are more likely to pursue sustained engagement in the same area, potentially reflecting greater absorptive capacity or strategic commitment. In contrast, firms entering unfamiliar domains may treat the acquisition as exploratory or one-off, with fewer follow-up investments.

¹³Figure A.7 in the appendix further explores the distribution of time between acquisition events by truncating the range to a maximum of six years. The shape of the distribution remains similar, with a large share of events clustered in the first two years. This robustness check supports our decision to model early follow-on acquisitions as part of a common post-treatment dynamic.

Figure 4: Big tech total patent share after acquisition




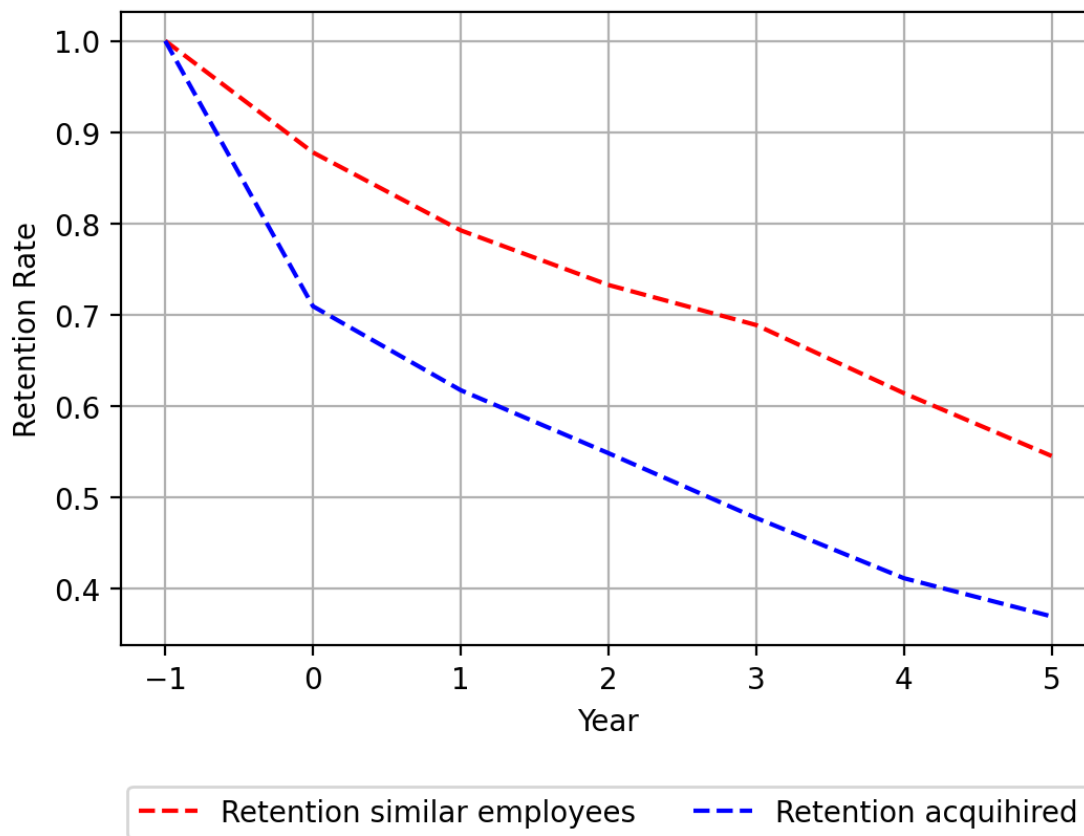
Figure 4 illustrates how acquisition waves are associated with concentration of patent ownership within a CPC. The figure plots the share of patents owned (collectively) by the eight large acquirers in the year after the k^{th} acquisition. We observe that the share of patents held by these big tech firms increases with cumulative acquisitions in a CPC group. This pattern reflects consolidation through two channels: patenting by the large acquirers, and acquisition of patents previously held by targets. The overall effect is that IP ownership is more concentrated in the hands of big tech for those CPC groups associated with a larger acquisitions wave.

3.4 Employee Retention

Our final set of descriptive statistic focus on employee retention. For this analysis, we construct an unbalanced panel of employee position-year observations in a five-year window around the acquisition event (from two years before through two years after).¹⁴ Job stints that begin and end in the same calendar year are excluded to reduce noise from short-term or uncertain employment records. In cases where a resume includes both a private-sector job and an academic or educational affiliation in the same year, we retain only the private-sector position. However, an individual may still hold multiple job positions across different firms

¹⁴To avoid double counting overlapping positions, we attribute the year of a transition to the new position rather than splitting the year.

Figure 5: Retention Rates: Targets vs. Acquirer Employees with Similar Tenure Profiles



in the same year if the data clearly indicate distinct roles.

Figure 5 plots employee retention rates for workers employed at the target firm in the year prior to acquisition. We compare these retention outcomes to a matched sample of employees from the acquiring firm, constructed to have similar tenure profiles at baseline. Specifically, for each target employee in our sample, we identify an employee from the acquirer who began working there in the same year the target employee started at their organization. We focus on employees that work for the acquirer directly and exclude any subsidiaries that were targets in past acquisitions. If multiple acquirer employees meet this criterion, we select one at random to serve as the control employee. This matching approach helps isolate differences in post-acquisition retention patterns that are attributable to the acquisition event itself, rather than to career stage or job tenure.

The horizontal axis shows years relative to the acquisition date, while the vertical axis measures the share of employees who remain at their original employer in each subsequent year. The figure reveals a stark divergence in retention dynamics following the acquisition. While retention among matched acquirer employees declines, retention among target firm

employees falls much more sharply starting in the year of acquisition and continues to fall in subsequent years.

Specifically, target employees experience a steep drop in retention in the first year after the deal, with almost 40% leaving their employer within a year. By three years after the acquisition, cumulative attrition exceeds 50%, suggesting significant turnover among the acquired workforce. In contrast, for most acquiring firms, the matched acquirer employees display a flatter retention curve, with substantially less attrition over the same time window.¹⁵

These patterns highlight the fragility of talent retention in the wake of technology (startup) acquisitions. Despite the strategic importance of acquiring human capital in technology deals, the evidence suggests that much of the target firm’s workforce departs after acquisition events.

4 Results

4.1 CPC Group Level Analysis

We first examine whether and how patenting activity within specific technological fields (defined by CPC groups) changes after an acquisition by a large digital firm. By leveraging the granularity of CPC-by-year data, we trace innovation dynamics around acquisition events and assess whether these transactions coincide with shifts in the rate of innovation.

4.1.1 Baseline Event Study

The goal of this section is to estimate how patenting activity within a technological domain evolves before and after acquisitions. By organizing the data at the CPC group-year level, we can track how innovation trends shift in response to a treatment event, which we define as an acquisition of a startup that had previously filed patents in that CPC group.

Our primary outcome variable, $lpat_{ct}$, is the arcsin transformation of the total number of patents filed in CPC class c in year t . Our baseline specification to evaluate the innovation

¹⁵Figure A.10 in the appendix shows that this retention pattern is broadly consistent across all eight acquiring firms in our sample. While there is some variation in the level and slope of post-acquisition attrition (with Apple exhibiting the highest overall retention) every acquirer displays a substantial decline in target employee retention following the acquisition. However, for two firms, Cisco and Intel, attrition seems similar for both (acquired) target employees and directly acquired employees.

trend in a technological area after a startup acquisition is as follows:

$$lpat_{ct} = \alpha + \gamma_c + \lambda_t + \sum_{s=-10}^{10} \tau_s \times d_{c,t=y+s} + \epsilon_{ct} \quad (1)$$

where γ_c are CPC group fixed effects and λ_t are calendar-year fixed effects. The indicator variable $d_{c,t=y+s}$ is set equal to one s years after a digital merger happened in calendar-year y in CPC category c . The coefficients τ_s – often called dynamic treatment effects – can be interpreted as the percentage change in patent filings relative to the baseline year ($s = -1$), controlling for both CPC group and calendar year fixed effects.

Figure 2 shows that we often have multiple acquisitions within a CPC group. When two or more acquisitions occur in a CPC group in the *same year*, we treat them as a single event at the CPC-year level. This approach helps to avoid double-counting and simplifies the structure of the event-study analysis while still capturing the cumulative impact of overlapping acquisitions. It is less clear how we should handle multiple acquisitions in the same CPC group that occur in different years. We begin by focusing on the first event, with follow-on acquisitions providing a potential source of heterogeneity in treatment effects. We then consider a specification, where we combine dummies for the first, second, and so on treatment events.

The implicit control group in equation 1 are the CPC groups without any startup acquisition event. Thus, our event study specification can be considered as a hybrid approach: identification comes from both variation in event times and the comparison to untreated CPC groups. In general, event study analyses have been found to be potentially biased, with robust estimators suggested by Callaway and Sant’Anna (2021) and Sun and Abraham (2021). Since our sample includes a large number of untreated CPC groups, we have found that our simpler specification yields near-identical results to the state-of-the-art estimators.¹⁶

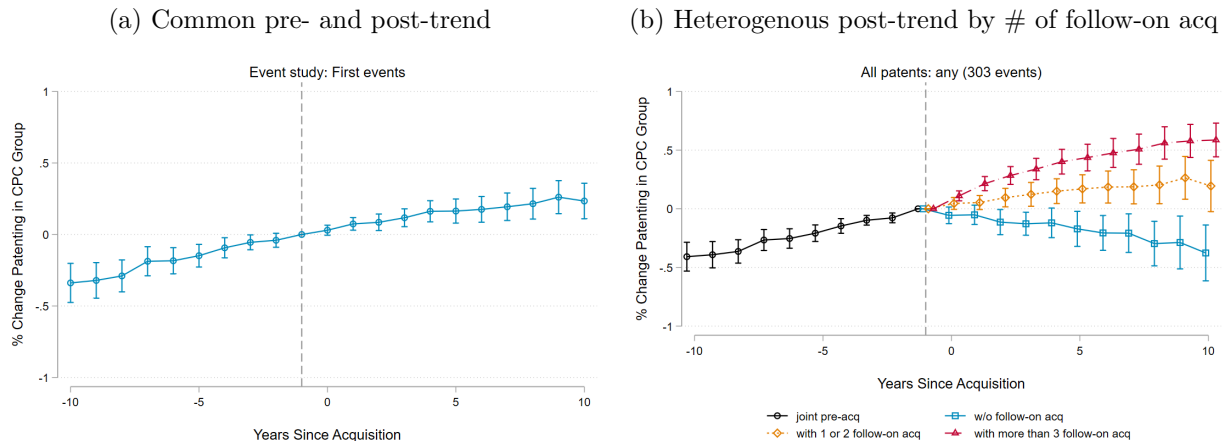
Figure 6 presents the estimated event-study coefficients τ_s for the effect of acquisitions on total patenting within treated CPC groups shown in equation (1).¹⁷ Panel (a) uses a pooled specification that considers the first acquisition in each CPC group as the treatment event and estimates a single set of pre- and post-treatment effects. The results reveal a marked upward trend in patenting before the acquisition, followed by continued growth afterward.

Panel (b) disaggregates the post-treatment effects by the number of follow-on acquisitions that occur in the same CPC group. This specification shows that the pre-acquisition growth in patent filings (relative to CPC codes without acquisitions) continues if and only

¹⁶It might be interesting to explore the differences in outcomes to a conventional event study with only treated units as mentioned by Miller (2023).

¹⁷Graphs by acquirer can be found in Appendix A.

Figure 6: Outcome: all patents within treated CPC codes, First acquisitions within CPC group



if there are follow-on acquisitions in the treated CPC group. When there are more follow-on acquisitions, the post-first-acquisition growth in patenting is larger. Patenting declines in treated CPC codes that see no further acquisition activity. This pattern suggests that innovation tends to rise most strongly in technological areas where digital incumbents continue to acquire patenting startups, possibly reflecting ongoing strategic investment in the underlying technology.

At the same time, both panels reveal strong pre-trends, with patenting activity already increasing in the years leading up to the first acquisition. This implies that acquisitions are not randomly assigned across CPC groups but are more likely to occur in areas with accelerating innovation. These pre-trends complicate a causal interpretation.

The presence of strong pre-trends in Figure 6 shows the existence of selection on observables: acquisitions tend to occur in CPC groups that are already experiencing rising innovation. To better isolate the effect of acquisitions, we next implement a matching approach that pairs treated CPC groups with observationally similar (but untreated) CPC groups that have comparable levels and trends of pre-treatment patenting. This allows us to construct a more plausible counterfactual and assess whether the observed post-acquisition patterns persist when comparing more closely aligned trajectories.

To be clear, there still might be selection on unobservable characteristics of the CPC groups. In particular, acquirers might have inside knowledge about future patenting activity in especially valuable CPC groups, leading them to acquire targets with technology in areas that they know (or expect) will grow in the future. Hence, finding an acceleration of patenting activity is not necessarily the results of a positive innovation effect of big tech acquisitions.

Table 6: Patent numbers and growth rate treated versus matched control CPC groups

	(1) Treated Mean	(2) Untreated Mean	(3) (2) – (1) b	(4) Matched Mean	(5) (4) – (1) b
Growth rate	0.09	-0.03	-0.11***	0.08	-0.01
Patent count by application date	255.75	27.17	-228.59***	246.25	-9.50
Observations	275	27040	27315	275	550

4.1.2 Event Study with Matched Controls

The preceding analysis relied on untreated CPC groups as a general control group. However, CPC groups that were never affected by an acquisition may differ systematically from those that were, especially in terms of pre-treatment innovation trends. To address this concern, we construct a more comparable control group using nearest-neighbor matching based on pre-acquisition patent growth, number of patent applications, and the acquisition year.

Specifically, we estimate the propensity score for experiencing a first acquisition event based on the growth rate in the year prior to the first acquisition and the total number of patent applications in the year of acquisition in the CPC group. For each treated CPC group, we identify the untreated CPC group with the closest propensity score and use it as a control. We then assign a pseudo-event year to the matched control that corresponds to the treated group’s acquisition year. When multiple treated CPC groups are matched to the same untreated group, the control observation is duplicated, and each copy inherits the respective event time. We retain only these matched pairs (i.e., treated and matched controls) in the estimation sample. Overall, this matching strategy helps ensure that treatment and control groups follow similar trends before the acquisition event, and are of similar size, allowing for a more credible comparison of post-acquisition outcomes.

Table 6 reports summary statistics for treated, unmatched untreated, and matched untreated CPC groups. The goal is to assess how well the matched control groups resemble treated groups along key pre-treatment dimensions. Column (1) shows the average patent growth rate in the year before acquisition and total number of patents for treated CPC groups. Column (2) presents the same statistics for all untreated CPC groups, while column (4) restricts the comparison to the matched subset. Columns (3) and (5) report the differences between treated and untreated groups for the unmatched and matched samples, respectively.

The unmatched untreated CPC groups exhibit substantially lower growth rates and much smaller patent counts than the treated groups, with large and statistically significant dif-

ferences. By contrast, the matched control groups closely resemble treated CPC groups in terms of pre-acquisition growth, with no significant difference in trends, and total application numbers in the year of acquisition. These comparisons confirm that the matching procedure effectively balances pre-treatment trends and helps to construct a more credible control group for the event study analysis.

After this matching procedure we run the following regression

$$lpat_{ct} = \alpha + \gamma_c + \lambda_t + \sum_{s=-10}^{10} \nu_s \times d_{t=y+s} + \sum_{s=-10}^{10} \tau_s \times d_{c,t=y+s} + \epsilon_{ct} \quad (2)$$


where γ_c are CPC group fixed effects and γ_t are year fixed effects. The indicator variable $d_{t=y+s}$ equals one for *both* the treated and matched control CPC groups s years after a digital merger happened (with ν_s as the corresponding coefficient). The indicator variable $d_{c,t=y+s}$ equals one for *only* the treated CPC category c . The coefficients of interest are the various τ_s . Each τ_s measures how patenting in treated CPC groups differs from patenting in matched controls s years after the acquisition event. Since we matched on pre-acquisition patent growth, there should be no significant difference in outcomes before the acquisition, at least for the five years on which we based our matching.

In line with our previous unmatched analysis of acquisition events, we investigate how innovation trends in treated and control CPC groups compare, differentiating by whether and how many follow-on acquisitions occur in a given CPC group. Figure 7 plots the estimated coefficients τ_s from the matched event study design. Panel (a) considers all first acquisition events and estimates a common post-treatment effect, assuming homogeneous responses across CPC groups. Panel (b) relaxes this assumption and estimates separate post-treatment trends based on the number of follow-on acquisitions observed in each group.

To ensure that matching is based on comparable pre-treatment dynamics, we restrict the sample to CPC groups with at least two years of data prior to the first acquisition. This requirement reduces the number of treated events to 275, excluding 28 events that occurred before the year 1997.

As expected, the use of matched control groups eliminates the strong pre-trends observed in the unmatched analysis, indicating improved balance between treated and control CPC groups. Post-acquisition, the results show clear heterogeneity: CPC groups with multiple follow-on acquisitions experience significantly larger increases in patenting. By contrast, CPC groups with only a single acquisition show no statistically significant change in the number of patents filed. This pattern reinforces the idea that follow-on acquisitions are associated with more persistent investment and innovation activity in the targeted technology

Figure 7: Matched event study; All first acquisition events




space.

The matched event study result in Figure 7 highlights the importance of accounting for the presence of further acquisition events when evaluating post-treatment innovation effects. To further understand whether the pronounced increase in patenting activity when there are follow-on acquisitions precedes or follows these further acquisitions, we run a similar regression as in equation 2, but using the *last acquisition event* within a CPC group as main treatment event. Figure 8 shows that, we do not observe an increase in patenting after the last acquisition event in treated CPC groups, compared to matched control CPC groups.¹⁸ This result provides further evidence that follow-on acquisitions are not merely coincidental but may reflect strategic investment in technological domains that are already experiencing innovation growth.

Overall, the evidence in this section shows that there is growth in patenting before the first startup acquisition in a CPC, and that this growth continues if *and only if* there are follow-on acquisitions in the same CPC. This pattern is consistent with models where innovation responds to *expectations* about exit or commercialization opportunities. It also highlights a potential limitation of our retrospective analysis: even if some acquisitions exert a strong influence on future innovation, the more important link between antitrust policy and innovation may operate through its ability to shape firms' ex-ante expectations. Finding sources of exogenous variation in those expectations will be challenging, however, in a lenient

¹⁸This is not a result of the matching procedure, as we observe the same qualitative patterns when running the unmatched event study with the last acquisition event as main treatment event. To further investigate the relative timing of innovation growth to the timing of acquisitions, we run a difference-in-difference specification that estimates the impact of successive acquisitions within a CPC group differentiating by presence of further acquisitions happening in the same group. The results can be found in Table A.5 in Appendix A.3.

Figure 8: Matched event study: All last acquisition events



environment where most startup acquisitions are approved.

While the analysis so far has focused on changes in the quantity of innovation (measured by patent application counts) at the CPC group level, it does not speak to the quality, influence, or diffusion of acquired patents. To address these dimensions, we now turn to citation-based measures, which offer a complementary perspective for evaluating the impact of acquisitions on technological significance and spillovers.

4.2 Patent Citation Analysis

In this section, we shift our focus (and unit of analysis) from the CPC group or technology area to the individual patent. Thus, instead of examining the change in overall innovation trends surrounding an acquisition event, we focus on the diffusion and impact of the acquired technology.

4.2.1 Impact of Acquisitions on Acquired Patents

Our main patent level outcome, cit_{pct} , is a count of forward citations to focal patent p , with primary CPC group c , in year t . Citations have been used extensively as a proxy for the technological significance, economic value, and cumulative innovation.

To examine the impact of acquisitions on the importance of acquired patents, we use a set of regressions that closely resemble those in [Rysman and Simcoe \(2008\)](#). Specifically, we analyze how the citation rate changes following acquisitions using the following models:

$$cit_{pct} = \alpha + \tau PostAcquisition_{pt} + Acquired_p + i.age_{pt} + \lambda_{ct} + \epsilon_{pct} \quad (3)$$

$$cit_{pct} = \alpha + \tau PostAcquisition_{pt} + age_{pt}^2 + age_{pt}^3 + \gamma_p + \lambda_{ct} + \epsilon_{pct} \quad (4)$$

$$cit_{pct} = \alpha + \sum_{s=-10}^{10} \tau_s \cdot d_{p,t=y+s} + age_{pt}^2 + age_{pt}^3 + \gamma_p + \lambda_{ct} + \epsilon_{pct} \quad (5)$$

The first specification of equation (3) is a pooled cross sectional model that includes an indicator, $Acquired_p$, for patents owned by the target firm; the diff-in-diff coefficient, $PostAcquisition_{pt}$, which equals one for acquired patents after the acquisition occurs; a full set of patent-age effects (where age is calculated relative to grant year) to capture the typical citation life-cycle; and a set of CPC-group by year fixed effects, λ_{ct} , that capture aggregate time-trends.

In equation (4), we introduce patent fixed effects γ_p , which absorb time-invariant differences in patent quality. Because these patent fixed effects also create a potential age-year-cohort co-linearity problem, however, we switch from using age dummies to a pair of terms, age_{pt}^2 and age_{pt}^3 , that capture the non-linear component of the underlying citation age profile. Finally, equation (5) implements an event-study specification, similar to our CPC-year analysis above, in order to estimate a full set of dynamic treatment effects τ_s .¹⁹

Table 7 presents the results from the baseline citation regressions described in specifications (3) and (4), using other patents in treated CPC groups as the control group.²⁰ Standard errors are clustered at the patent level. Column (1) corresponds to a specification with CPC-year fixed effects and a treatment indicator $Acquired_p$, while column (2) replaces the treatment dummy with patent fixed effects to account for time-invariant differences in patent quality and visibility. Both specifications include controls for patent age.

The key variable of interest, $PostAcquisition_{pt}$, is a dummy equal to one in all years after the acquisition of the patent's parent firm. The coefficient on this variable is positive and highly significant in both specifications, indicating that patents filed by acquired firms receive significantly more citations in the years following the acquisition. In column (1), the

¹⁹To clarify the structure of our citation-level analysis, Section A.1 in the appendix shows and discusses the distribution of treated patents by event year relative to acquisition. In particular, it highlights how our identification relies on a balanced panel of patent-year observations concentrated around the acquisition date.

²⁰In Section A.2 in the appendix, we report robustness checks using alternative control groups, including all patents in G and H, patents in untreated CPC groups, and patents assigned to the acquirer.

Table 7: Diff-in-Diff Estimates of Acquisition Impact on Patent Cites

	(1)	(2)	(3)	(4)	(5)	(6)
	#Citations	#Citations	#Citations	#Citations	#NonSelfCit	#PatFamCit
PostAcquisition	1.10*** (0.10)	1.13*** (0.09)	0.80*** (0.10)	0.89*** (0.10)	0.85*** (0.08)	0.43*** (0.06)
Acquired	0.54*** (0.05)		0.52*** (0.05)			
Post \times FirstNoFollowon			0.43 (0.45)	0.56 (0.42)		
Post \times FirstFollowon			2.15*** (0.58)	1.71** (0.54)		
CPC \times Year FE	yes	yes	yes	yes	yes	yes
Patent ID FE	no	yes	no	yes	yes	yes
Age FE	yes	no	yes	no	no	no
Age ² , Age ³	no	yes	no	yes	yes	yes
Adj. R ²	0.05	0.42	0.05	0.42	0.42	0.32
Avg. Outcome	0.80	0.80	0.80	0.80	0.80	0.48
N	44250965	44250965	44250965	44250965	44250965	44250965

Standard errors in parentheses. Column (3) includes main effects of First and Followon.


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

post-acquisition effect is approximately 1.10 additional citations per year, which corresponds to a 138 percent increase in the mean citation rate. The coefficient on $Acquired_p$ in column (1) indicates that the patents of acquired startups have a 67 percent higher citation rate than other patents in the CPC classes *before* the acquisition. This suggests that acquired patents are both more important to begin with, and see a large boost in their visibility and significance following an acquisition. In column (2), we add patent fixed effects, and find that the main treatment effect remains essentially unchanged.

The results in Table 7 indicate that acquisitions are followed by a marked increase in the visibility and influence of the acquired firm’s patents. Since the comparison group is composed of other patents in the same CPC groups, these results are not driven by broad shifts in technological fields, but rather reflect changes specific to the individual inventions of the acquired targets. This pattern is consistent with improved dissemination, commercialization, or reuse of the acquired knowledge, and suggests that acquisitions by digital incumbents may facilitate rather than suppress post-merger innovation spillovers.

Figure 9 depicts the estimated τ_s coefficients from the event-study specification (5). The specification includes patent fixed effects, CPC-year fixed effects, and controls for patent age (in quadratic and cubic form), and uses other patents in treated CPC groups as the control group. The event time s indexes years relative to the acquisition of the patent’s parent firm, with $s = 0$ corresponding to the acquisition year.

Figure 9: Event study: Application versus Grant Year Cites



The difference between the two panels in Figure 9 is driven by an assumption about citation timing. Panel (a) assigns forward citations to the application-year of the citing patent, and panel (b) assigns citations to the year when the citing patent is granted. The convention in the innovation literature is to assign cites to the application-year, based on the idea that this approximates the moment when an inventor utilized her knowledge of the prior art. This inevitably creates some measurement error, however, because citations can be added (by either the applicant or the examiner) during the examination process. If these “in process” citations are caused by an intervening event — such as the acquisition of the firm that filed the cited patent — it may be more appropriate to assign citations to the grant-year, as we do in panel (b).

In both panels of Figure 9, we observe a clear and sustained increase in citations to acquired patents following the acquisition. The magnitude of this effect is approximately one additional citation per year relative to pre-acquisition levels, or a doubling of the baseline citation rate. This is broadly consistent with the difference-in-differences results in Table 7, and supports the idea that acquisitions by large digital firms may amplify the diffusion of acquired innovations, rather than curtailing their impact.

While both panels in Figure 9 show an increase in the citation rate around the year of acquisition, the pattern of the dynamic treatment effects differs under the alternative timing assumptions. In panel (a), when forward cites are assigned to the citing patent’s application-year, we observe an increase in the citation rate during the five-year period before the acquisition, followed by a sharp acceleration in the year of acquisition, and then a leveling-off at a higher rate over the next decade. In panel (b), when forward cites are assigned to the grant-year, the difference in pre-acquisition trends disappears for the three-years before acquisition, and the dynamic treatment effects exhibit a larger and more sustained increase

over the post-acquisition years.²¹

The difference between the two graphs is consistent with the idea that the differential “pre trend” in panel (a) is caused by citations that arrive during the prosecution process (potentially as a consequence of the acquisition itself). Thus, we interpret the timing of the citation increase in panel (a), along with the disappearance of the pre-trend in Panel (b), as evidence that supports a causal interpretation of our diff-in-diff estimates. Of course, a more skeptical reader might treat them, instead, as descriptive estimates that capture both selection and treatment effects. Going forward, we follow the convention in the literature and assign citations to the application year of the citing patent.

4.2.2 Patent Citation Heterogeneity by Acquisition Sequence

While the average increase in citations suggests that acquisitions can enhance the impact of acquired innovations, this effect may vary depending on the context and sequence of acquisitions. In particular, we explore whether citation dynamics differ between first-acquired targets within a CPC group and those acquired later, and whether the presence of follow-on acquisitions amplifies or moderates these effects. This analysis helps to uncover whether the timing and positioning of an acquisition within a broader strategy shapes the extent to which innovation spillovers materialize.


Columns (3) and (4) in Table 7 present estimates from specifications (3) and (4), where we examine whether the citation effects of acquisitions differ depending on whether the patent belongs to the first acquired target within a CPC group or to a later target. The sample is split accordingly, and we further distinguish between cases where there were follow-on acquisitions and those where there were none. The control group in all regressions remains all other patents in treated CPC groups.

The interaction term *Post×FirstNoFollowon* captures the effect of an acquisition on patent citations for the first acquired target within a CPC group, conditional on there being no subsequent acquisitions in that group. Similarly, *Post×FirstFollowon* measures the post-acquisition citation effect for patents of the first target when there is a follow-on acquisition within the same CPC group. Finally, the coefficient on *PostAcquisition* represents the effect for patents associated with later acquisition events, that is, targets acquired after an earlier acquisition has already occurred in the CPC group.

The results reveal notable heterogeneity. Patents owned by first-acquired targets in CPC groups that subsequently experience additional acquisitions show a particularly strong post-acquisition increase in citations. This suggests that early movers in acquisition waves

²¹Panel (b) still shows a differential trend over the longer 10-year pre-acquisition window, but it is important to recognize that the composition of the sample is also shifting over that longer period (see Figure A.1).

Figure 10: Patents of first targets within CPC group




may benefit from greater organizational attention or resource integration, which in turn enhances the diffusion of their innovations. In contrast, patents associated with isolated acquisitions (i.e., those not followed by further deals) show smaller or insignificant citation effects, consistent with a more limited post-merger innovation impact. Later-acquired targets also exhibit positive post-acquisition citation effects, though generally smaller in magnitude compared to first-acquired firms in multi-acquisition CPC groups.

These findings suggest that the strategic sequencing of acquisitions matters for innovation outcomes. First acquisitions in an area that becomes the focus of sustained follow-on activity appear to generate the strongest citation-based spillovers.

Figures 10 and 11 display the estimated τ_s coefficients from the event-study specification (5), separately for patents of first-acquired and later-acquired targets within a CPC group. These figures provide a dynamic view of how the timing of acquisition within a technological domain affects the post-merger trajectory of innovation impact, as measured by forward citations.

Figure 10 focuses on patents belonging to the first target acquired within each CPC group. The event-study profile shows little movement in citations prior to the acquisition, followed by a pronounced and persistent increase beginning shortly after the acquisition year. The sharp and sustained post-treatment rise in citations suggests that first-acquired targets often experience a meaningful boost in the visibility and influence of their innovations. This

Figure 11: Patents of later targets within CPC group



effect is particularly strong in CPC groups that go on to experience follow-on acquisitions.

In contrast, Figure 11 plots the citation dynamics for patents of later-acquired targets. Here, the post-acquisition increase in citations is more muted and less sustained. While there is some upward movement following the acquisition year, the magnitude of the effect is smaller, and the confidence intervals are wider, especially in later years. These patterns suggest that later acquisitions may receive less organizational focus or integration effort, or may be motivated by defensive rather than developmental aims.


4.2.3 Mechanisms for the Citation Impact of Acquisitions

The preceding analysis documents a substantial post-acquisition increase in citations to the acquired firm's patents. But who is responsible for these additional citations? Do they reflect broader knowledge diffusion and technological spillovers, or are they primarily driven by increased internal use within the acquiring firm?

Figure 12 sheds light on these questions by showing how the share of total citations attributable to the acquiring firm evolves around the acquisition event. Specifically, it plots the fraction of all forward citations that originate from patents assigned to the acquirer, relative to the event year. The figure reveals a sharp rise in self-citations beginning immediately after the acquisition, suggesting that acquiring firms intensify their reuse or integration of acquired technologies. This internal uptake likely reflects strategic alignment or absorptive

capacity following the merger.

Figure 12: Patents of later targets within CPC group




To assess whether the overall citation increase is limited to the acquirer or reflects broader innovation spillovers, we re-estimate the event-study specification from earlier sections, but exclude all citations made by the acquiring firm. Figure 13 displays the resulting event-study coefficients, focusing exclusively on citations from non-acquirer patents.

The figure shows that while the magnitude of the effect is somewhat reduced relative to the full-sample specification (see Figure 9), a strong and sustained post-acquisition increase in citations remains. Importantly, the pre-trend flattens further in this specification, reinforcing the idea that internal citations may partly drive early citation growth. Nonetheless, the continued post-acquisition rise in non-acquirer citations suggests that the acquired knowledge becomes more visible and influential in the broader technological ecosystem and not just within the acquiring firm.

Column (5) of Table 7 reports the corresponding difference-in-differences estimates for the non-acquirer citation sample. The results are consistent with the event-study graph: the coefficient on the post-acquisition indicator remains positive and statistically significant, although somewhat smaller in magnitude. The estimated effect is approximately 0.85 citations per year, down from 1.13 in the baseline specification that includes all citations. This confirms that the observed citation effects are not exclusively driven by self-citations and reflect genuine external spillovers.

Figure 13: Non-acquirer citations only



These results provide compelling evidence that the citation boost observed after acquisitions is not simply an artifact of internal referencing by the acquirer. Rather, acquired patents experience greater visibility and influence within the broader innovation ecosystem, with more citations coming from unaffiliated firms. The fact that the effect persists after removing self-citations underscores the potential for knowledge spillovers generated by integration, diffusion, or signaling effects following acquisition. Taken together with the event-study results in Figure 13, column (5) of Table 7 strengthens the case that acquisitions can enhance rather than hinder the external relevance of the acquired firm’s technology.

Thus, acquisitions by large digital firms appear to enhance both internal reuse and external diffusion of acquired technologies. While acquirers themselves are often the first to benefit from newly integrated capabilities, the effects extend beyond the boundaries of the firm and contribute to broader innovation dynamics.

In addition to distinguishing self-cites from external cites, we might want to aggregate all citations from “families” where many patents correspond to the same underlying invention. To implement that idea, we aggregate citations across all patents belonging to the same DOCDB family as the focal patent. This approach accounts for cases where the same invention is patented in multiple jurisdictions or where new claims are added through continuation applications, potentially leading to “duplicative” citations (Kuhn, Younge and Marco, 2020).

Column (6) of Table 7 and Figure A.5 in the appendix report results using citations

aggregated at the patent family level. The pattern is again consistent with our previous analysis. There is little to no pre-trend, followed by a clear and sustained post-acquisition rise in citations. The estimated effect is approximately 0.43 citations per year which is still economically and statistically significant though substantially smaller than our baseline results.

These alternative specifications reinforce the core finding that acquisitions are associated with an increase in the visibility and diffusion of acquired patents. The robustness of the results to patent family aggregation suggests that our main conclusions are not driven by mechanical features of patent structure.

4.3 Employee Retention and Citations

In this section, we use workforce data from Revelio Labs to examine how employee retention following acquisitions shapes the diffusion of acquired technology. By focusing on employee tenure and turnover, we aim to understand how the stability of the acquired workforce influences the diffusion of innovation as measured by citations.

To explore the link between post-acquisition workforce integration and the citation impact of acquired patents, we examine whether citation effects vary with employee retention. As Section 3.4 shows, retention of target firm employees drops sharply in the years following an acquisition, but the extent of this decline varies across deals. If innovation outcomes depend not only on the acquisition itself but also on the successful integration of human capital, we should expect heterogeneity in citation effects based on how many employees remain at the acquiring firm. We test this hypothesis by interacting the post-acquisition treatment indicator with the (demeaned) share of target employees still employed at the acquirer three years after the deal.

Table 8 presents estimates from a regression specification that examines whether the citation effects of acquisitions vary with post-acquisition employee retention. Specifically, we interact the post-treatment indicator with the (demeaned) share of target employees still employed at the acquiring firm three years after the acquisition, using data from Revelio Labs (as shown in Figure 5). The dependent variable remains annual forward citations per patent, and we include the same fixed effects and age controls as in prior specifications.

Table 8 presents regression results examining how the post-acquisition citation effect varies with the share of the target workforce retained at the acquiring firm. Column (1) includes CPC group-year and age fixed effects, while column (2) uses patent fixed effects and higher-order polynomial age controls instead of age fixed effects.

Consistent with our previous analysis, the coefficient on the post-acquisition indicator

Table 8: Citation effect by retention rate at the acquirer

	(1)	(2)
	#Citations	#Citations
PostAcquisition	0.79*** (0.09)	0.91*** (0.09)
Post \times Retention	-0.97** (0.38)	-0.68* (0.36)
Acquired	0.46*** (0.06)	
Acquired \times Retention	0.40* (0.23)	
CPC \times Year FE	yes	yes
Patent ID FE	no	yes
Age FE	yes	no
Age ² , Age ³	no	yes
Adj. R2	0.05	0.42
Avg. Outcome	0.82	0.82
SD Retention Rate	0.19	0.19
N	41878689	41878689

Clustered standard errors in parentheses

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

is positive and significant in both columns: 0.79 in column (1) and 0.91 in column (2), confirming that acquired patents receive more citations after the acquisition. However, the interaction between the post-acquisition indicator and the demeaned retention rate is negative and statistically significant in both specifications (-0.97 and -0.68 , respectively). These coefficients imply that the citation boost from acquisition is smaller in deals with higher employee retention. At the sample mean retention rate of 45%, the estimated citation gain is near the baseline level, but it declines substantially as retention increases. For instance, based on column (2), acquisitions with the highest retention rates in our sample (around 85%) experience a citation gain that is lower by 0.3 citations per year compared to acquisitions with mean retention rates.

Because column (1) does not include patent fixed effects, we also include a dummy for treated patents (those belonging to acquired firms) as well as an interaction between this dummy and the demeaned retention rate. The positive coefficient on the treated dummy (0.46) indicates that acquired patents already had more citations prior to the acquisition compared to patents from non-acquired firms. Moreover, the positive and statistically sig-

nificant interaction term (0.4) shows that this pre-treatment citation advantage was even greater for patents from targets that would go on to have higher employee retention. This reinforces the interpretation that baseline differences in citation levels were associated with retention-related characteristics of the targets. At the same time, the post-treatment interaction remains negative and significant, suggesting that although high-retention targets started with a higher citation baseline, their acquired patents experienced a smaller increase in citations after the acquisition. This pattern is consistent with the idea that innovation spillovers may be more pronounced when employees leave the acquirer, diffusing knowledge into the broader ecosystem.

Taken together, these results suggest that greater retention of target employees may dampen the citation surge typically observed after acquisitions. or, in more colloquial terms, a successful acqui-hire may be bad news for spillovers. One plausible explanation is that when retention is low, departing employees may disseminate knowledge externally by joining or founding other firms, leading to broader diffusion and more citations. This mechanism is consistent with the model of Silicon Valley “job hopping” proposed in [Fallick, Fleischman and Rebitzer \(2006\)](#). Conversely, when more of the workforce is retained, innovations may remain embedded within the acquiring firm, limiting external visibility. The results are robust across both specifications and offer new insight into the relationship between labor mobility and post-merger knowledge spillovers.

In addition, these results suggest that citation effects following acquisition are not driven by continuity of the target workforce. Instead, they may be amplified in cases where acquisitions trigger more substantial organizational restructuring, redirection of innovation resources, or knowledge spillovers through employee mobility. Thus, employee departures may not only signal a breakdown in integration, but may also facilitate broader dissemination of innovation beyond the acquiring firm’s boundaries.

5 Conclusion

In this paper we investigated the impact of acquisitions by major digital incumbents on innovation in targeted technological domains. Using a novel dataset that links over 1,000 startup acquisitions by eight large technology firms to detailed patent-level information, we implemented an event-study framework to trace changes in both the quantity and quality of innovation before and after acquisition events. By analyzing patenting trends within Cooperative Patent Classification (CPC) groups, and tracking citation patterns to acquired patents, we provided a comprehensive assessment of how these deals shape the evolution of technological activity.

Our results offer a nuanced view of the “killer acquisition” narrative. While acquisitions by large digital firms are often viewed with suspicion, we find little evidence that these deals suppress innovation. On the contrary, patenting activity tends to increase in CPC groups following an acquisition, especially when there are follow-on acquisitions in the same domain. Matched control group comparisons confirm that this pattern is not driven by differential pre-trends. Citation-based analyses further reinforce these findings: patents owned by acquired firms receive significantly more citations after the acquisition, and this increase is not solely due to self-citations by the acquirer. Spillovers to external innovators play a meaningful role, suggesting broader knowledge diffusion rather than internal hoarding.

Importantly, we document heterogeneity in these effects. The strongest post-acquisition innovation responses occur when the acquirer had prior patenting experience in the CPC group and when the acquisition is followed by additional deals in the same domain. First acquisitions in these acquisition waves show especially large increases in citations, indicating that the strategic sequence and technological alignment of acquisitions matter for post-merger innovation outcomes. Robustness checks that use grant-year citations and patent family aggregations confirm the reliability of our results.

From a policy perspective, our findings suggest that many acquisitions by digital incumbents are motivated by complementarity rather than suppression. While this does not rule out the existence of killer acquisitions in specific cases, the evidence indicates that, on average, these deals promote rather than stifle technological progress. These insights have implications for merger review in digital markets, where blanket skepticism may overlook the potential for innovation-enhancing firm integration.

References

- Akcigit, Ufuk, John Grigsby, and Tom Nicholas**, “Innovation and Top Income Inequality,” *The Review of Economic Studies*, 2016, 85 (1), 1–45.
- Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella**, “Markets for Technology and Their Implications for Corporate Strategy,” *Industrial and Corporate Change*, 2004, 13 (6), 935–953.
- Arora, Asish and Alfonso Gambardella**, “The Market for Technology,” in Bronyn Hall and Nathan Rosenberg, eds., *Handbook for the Economics of Innovation*, Vol. 1, Elsevier, 2010, pp. 641–678.
- Asil, Aslihan, Paulo Ramos, Amanda Starc, and Thomas G Wollmann**, “Painful Bargaining: Evidence from Anesthesia Rollups,” *NBER Working Paper*, 2024.

- Bloom, Nicholas, Mark Schankerman, and John Van Reenen**, “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica*, 2013, *81* (4), 1347–1393.
- Bryan, Kevin A and Erik Hovenkamp**, “Antitrust limits on startup acquisitions,” *Review of Industrial Organization*, 2020, *56* (4), 615–636.
- Buehler, Stefan and Armin Schmutzler**, “Mergers and Innovation: The Role of R&D Spillovers,” *RAND Journal of Economics*, 2005, *36* (4), 995–1012.
- Callaway, Brantly and Pedro HC Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Cunningham, Colleen, Florian Ederer, and Song Ma**, “Killer Acquisitions,” *Journal of Political Economy*, 2021, *129* (3), 649–702.
- Fallick, Bruce, Charles A. Fleischman, and James B. Rebitzer**, “Job-Hopping in Silicon Valley: Some Evidence concerning the Microfoundations of a High-Technology Cluster,” *The Review of Economics and Statistics*, 2006, *88* (3), 472–481.
- Federal Trade Commission**, “Non-HSR Reported Acquisitions by Select Technology Platforms, 2010-2019,” 2021. <https://www.ftc.gov/system/files/documents/reports/non-hsr-reported-acquisitions-select-technology-platforms-2010-2019-ftc-study/p201201technologyplatformstudy2021.pdf>.
- Federico, Giulio, Fiona Scott Morton, and Carl Shapiro**, “Antitrust and innovation: Welcoming and protecting disruption,” *Innovation Policy and the Economy*, 2020, *20* (1), 125–190.
- Federico, Giulio, Gregor Langus, and Tommaso Valletti**, “A simple model of mergers and innovation,” *Economics Letters*, 2017, *157*, 136–140.
- Federico, Giulio, Gregor Langus, and Tommaso Valletti**, “Horizontal mergers and product innovation,” *International Journal of Industrial Organization*, 2018, *59*, 1–23.
- Fumagalli, Chiara, Massimo Motta, and Emanuele Tarantino**, “Shelving or developing? Optimal policy for mergers with potential competitors,” *CEPR Discussion Paper*, 2022.
- Gaffard, Jean-Luc and Francesco Quatraro**, “Merger Policy and Innovation: The Role of Product Market Competition,” *Industrial and Corporate Change*, 2022, *31* (5), 1047–1076.
- Gans, Joshua S. and Scott Stern**, “Incumbency and R&D Incentives: Licensing the Gale of Creative Destruction,” *Journal of Economics & Management Strategy*, 2000, *9* (4), 485–511.
- Gans, Joshua S., David H. Hsu, and Scott Stern**, “When does start-up innovation spur the gale of creative destruction?,” *RAND Journal of Economics*, 2002, *33* (4), 571–586.

- Gawer, Annabelle, Michael A Cusumano et al.**, *Platform leadership: How Intel, Microsoft, and Cisco drive industry innovation*, Vol. 5, Harvard Business School Press, Boston, 2002.
- Gilbert, Richard**, *Innovation Matters: Competition policy for the high-technology economy*, MIT press, 2020.
- Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg**, “The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools,” NBER Working Paper No. 8498, National Bureau of Economic Research 2001.
- Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson**, “Geographic localization of knowledge spillovers as evidenced by patent citations,” *Quarterly Journal of Economics*, 1993, pp. 577–598.
- Jin, Ginger Zhe, Mario Leccese, and Liad Wagman**, “How do top acquirers compare in technology mergers? New evidence from an S&P taxonomy,” *International Journal of Industrial Organization*, 2023, 89, 102891.
- Kamepalli, Sai Krishna, Raghuram Rajan, and Luigi Zingales**, “Kill zone,” Technical Report, National Bureau of Economic Research 2020.
- Kuhn, Jeffry, Kenneth Younge, and Alan Marco**, “Patent citations reexamined,” *The RAND Journal of Economics*, 2020, 51 (1), 109–132.
- Lemos, Thales de Melo e and Guilherme Mendes Resende**, “Nascent Competition and Killer Acquisitions in Digital Markets: A Decade of Acquisitions by Big Tech,” *ResearchGate Preprint*, 2023.
- Letina, Igor, Armin Schmutzler, and Regina Seibel**, “Killer Acquisitions and Beyond: Policy Effects on Innovation Strategies,” *International Economic Review*, 2024, 65 (2), 591–622.
- Miller, Douglas L.**, “An Introductory Guide to Event Study Models,” *Journal of Economic Perspectives*, 2023, 37 (2), 203–230.
- Motta, Massimo and Sandro Shelegia**, “The “Kill Zone”: When a Platform Copies to Eliminate a Potential Threat,” *Journal of Economics & Management Strategy*, 2025, 34 (3), 657–673.
- OECD**, “Start-ups, Killer Acquisitions and Merger Control,” Technical Report, Organisation for Economic Co-operation and Development 2020.
- Olmsted-Rumsey, Jane, Federico Puglisi, and Liangjie Wu**, “Digital Platform Acquisitions and Growth,” *Working Paper*, 2024.
- Poege, Felix**, “Competition and Innovation: The Breakup of IG Farben,” *Workgin Paper*, 2024.


- Rysman, Marc and Timothy Simcoe**, “Patents and the Performance of Voluntary Standard-Setting Organizations,” *Management Science*, 2008, *54* (11), 1920–1934.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.
- Teece, David J.**, “Profiting from Technological Innovation: Implications for Integration, Collaboration, Licensing And Public Policy,” *Research Policy*, 1986, *15* (6), 285–305.
- U.S. Department of Justice and Federal Trade Commission**, “Merger Guidelines,” 2023. https://www.ftc.gov/system/files/ftc_gov/pdf/2023_merger_guidelines_final_12.18.2023.pdf.
- Valletti, Tommaso**, “The innovation theory of harm in merger control: Some clarifications,” *Economics Letters*, 2025, p. 112556.
- Valletti, Tommaso M. and Hans Zenger**, “Merger control and innovation,” *Journal of Competition Law and Economics*, 2017, *13* (3), 349–370.
- Watzinger, Martin, Thomas Fackler, Markus Nagler, and Monika Schnitzer**, “How Antitrust Enforcement Can Spur Innovation: Bell Labs and the 1956 Consent Decree,” *American Economic Journal: Economic Policy*, 2020, *12* (4), 328–359.

A Appendix

A.1 Citation Analysis Structure

To better understand the structure of the citation analysis, Figure A.1 plots the number of treated patents by event time (i.e., by the number of years before or after the acquisition of their parent firm). This figure provides insight into the distribution of observation counts across event years in the citation-level regressions. The horizontal axis indicates the number of years relative to the acquisition event (with zero denoting the year of acquisition), while the vertical axis shows the total number of patent-year observations in each relative year.

Figure A.1: Number of treated patents by year relative to acquisition



The distribution is roughly symmetric around the acquisition year, but with a notable concentration of observations in the period spanning five years before to five years after acquisition. This reflects both the temporal clustering of patent filings around the acquisition event and our balanced panel construction, which restricts some specifications to patents with sufficient data on both sides of the event window. The drop-off in observations beyond these bounds is due to right- and left-censoring in the patent dataset.

Importantly, this pattern provides some context for the event-study estimates presented later in the citation analysis. Since the majority of variation is concentrated within a ten-year window surrounding acquisition, the identification of pre- and post-trends relies most heavily on this central range. The balanced coverage across years also supports the credibility

of dynamic citation profiles, allowing us to assess both anticipatory trends and longer-run post-acquisition effects.

A.2 Different Control Groups

In the main citation analysis, we use other patents in treated CPC groups as the control group. This ensures that the treated and control patents are exposed to the same broad technological environment, but it also means that any citation spillovers within the CPC group are captured as part of the treatment effect. To assess the robustness of our results to alternative definitions of the control group, we consider three complementary specifications: (1) all patents in CPC sections G and H, (2) patents in untreated CPC groups, and (3) other patents owned by the acquiring firm.

Each specification addresses a distinct concern. The G&H sample allows us to test whether the observed citation effects persist when comparing to a much broader cross-section of digital technologies. The untreated CPC group sample restricts attention to patents in domains untouched by acquisitions, offering a cleaner comparison free from potential intra-group spillovers. Finally, using other patents by the acquirer allows us to ask whether the citation gains of the acquired firm’s patents exceed those of the acquirer’s own preexisting technologies.

A.2.1 All Patents in G and H

As a first robustness check, we broaden the control group to include all patents in CPC sections G (Physics) and H (Electricity). These sections cover a wide array of digital technologies, including semiconductors, computing, telecommunications, and electronic control systems. By comparing acquired patents to this broader technological baseline, we can assess whether the observed citation effects are specific to narrowly defined CPC groups or reflect more general changes in patent influence.

Table A.1 reports regression results using this expanded control group for specifications (3) and (4). The estimates are similar in magnitude and significance to our main results. In column (1), which includes CPC-year fixed effects and a treatment indicator, the post-acquisition effect is 1.17 citations per year. In column (2), which adds patent fixed effects and polynomial age controls, the effect remains strong at 1.16 citations per year. These coefficients confirm that acquired patents receive substantially more citations following the acquisition, even relative to a much broader cross-section of patents in related technological fields.

Figure A.2 presents the corresponding event-study coefficients from specification (5). The

Table A.1: Control group: All patents in G and H

	(1)	(2)
	#Citations	#Citations
Post	1.17*** (0.10)	1.16*** (0.09)
Treated	0.50*** (0.05)	
Age ²		-0.02*** (0.00)
Age ³		0.00*** (0.00)
CPC \times Year FE	yes	yes
Patent ID FE	no	yes
Age FE	yes	no
Adj. R2	0.05	0.41
Avg. Outcome	0.73	0.73
N	61916941	61916937

Standard errors in parentheses

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

graph shows a flat pre-trend, followed by a pronounced and sustained increase in citations beginning in the acquisition year. This pattern closely mirrors the baseline event study shown in Figure 9, indicating that the citation gains are not driven by idiosyncratic comparisons within specific CPC groups. Rather, the boost in patent visibility appears to be a robust and generalizable feature of the post-acquisition environment.

Overall, these results demonstrate that our findings are not sensitive to the choice of control group. Even when using all G and H patents which include many unrelated technologies, as a benchmark, acquired patents stand out in terms of increased citation activity.

A.2.2 Patents in Untreated CPC Groups

As a second robustness check, we restrict the control group to patents in CPC groups that are never affected by any acquisition in our sample period. This approach offers a cleaner comparison group, free from potential intra-group spillovers that might contaminate estimates when using patents from treated CPC groups as controls. However, this control group may also differ systematically in technological intensity or market relevance, which we address through fixed effects and specification controls.

Figure A.2: Event study, control group: All patents in G and H

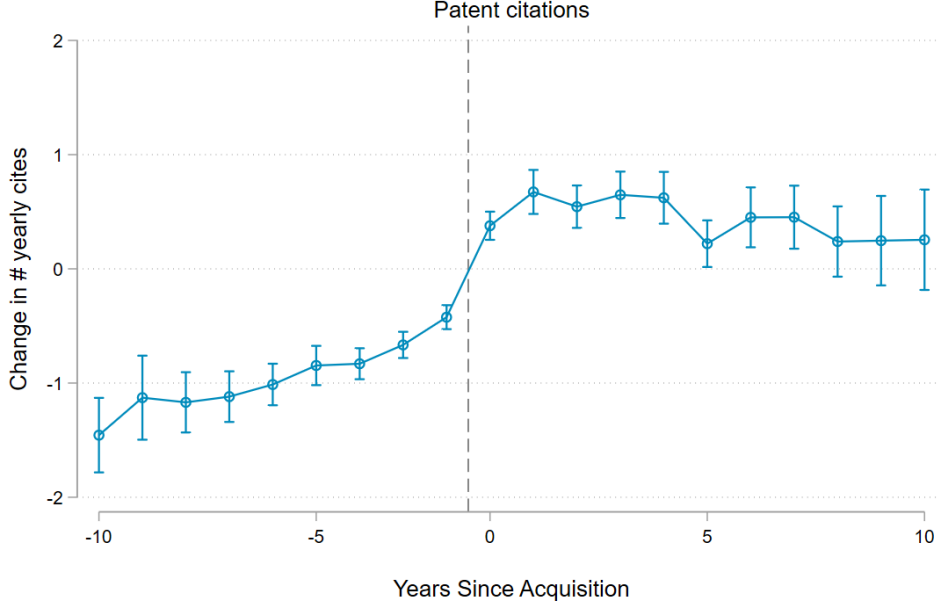


Table A.2 presents the regression results using untreated CPC groups as the comparison set for specifications (3) and (4). In both specifications, the coefficient on the post-acquisition indicator remains large and statistically significant. In column (1), the estimated effect is 2.20 citations per year which is substantially higher than in the baseline. In column (2), which includes patent fixed effects and polynomial age controls, the estimate is 1.71. The larger magnitude of these effects is likely due to differences in baseline citation levels: the untreated CPC groups tend to be less active, with an average citation rate of 0.55 compared to 0.80 in the treated sample. As a result, any post-treatment shift appears relatively larger in absolute terms.

Figure A.3 shows the event-study coefficients τ_s from specification (5), again using untreated CPC groups as the control. The dynamic pattern is consistent with our core findings: citations to acquired patents begin to rise sharply in the year of acquisition and continue to grow over the following years. Importantly, there is no evidence of a pre-trend, further supporting the credibility of the event-study design. While the magnitude of the post-treatment rise is slightly higher in this specification, the overall trajectory closely matches the baseline results.

These findings provide further support for the interpretation that acquisitions lead to increased technological influence of the acquired patents. Even when comparing to entirely unaffected CPC domains, the observed rise in citations persists.

Table A.2: Control group: Patents in untreated CPC groups

	(1)	(2)
	#Citations	#Citations
Post	2.20*** (0.21)	1.22*** (0.09)
Age ²		-0.02*** (0.00)
Age ³		0.00*** (0.00)
CPC \times Year FE	yes	no
Year FE	no	yes
Patent ID FE	no	yes
Age FE	yes	no
Adj. R2	0.05	0.34
Avg. Outcome	0.55	0.55
N	17740111	17744080

Standard errors in parentheses

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


A.2.3 Other Patents by Acquirer

As a final robustness check, we use as the control group all other patents owned by the acquiring firm. This comparison allows us to test whether the post-acquisition increase in citations to acquired patents reflects general firm-level dynamics such as an overall rise in patent visibility, increased marketing, or changes in citation practices within large firms or whether the acquired patents exhibit a uniquely strong response. If citation growth is observed only for acquired patents and not for the acquirer's own portfolio, this strengthens the case that the effect is driven by the acquisition itself rather than broader shifts at the firm level.

Table A.3 reports the corresponding regression estimates for specifications (3) and (4). In column (1), which includes CPC-year fixed effects and a treatment indicator, the post-acquisition coefficient is 0.92 citations per year. In column (2), which adds patent fixed effects and flexible age controls, the estimate is slightly lower at 0.88 citations. While these magnitudes are smaller than in the baseline specification using other treated CPC patents as the control group, they remain statistically and economically significant. Importantly, the average citation rate in the acquirer's patent portfolio is relatively high (0.98), making this a conservative benchmark.

Figure A.4 shows the event-study coefficients from specification (5), again comparing

Figure A.3: Event study, Patents in untreated CPC groups



acquired patents to other patents owned by the acquiring firm. The pre-trend is relatively flat, indicating no divergence in citation trajectories prior to the acquisition. After the acquisition, however, citations to acquired patents increase noticeably relative to the firm's own baseline. This suggests that the integration of external technologies delivers a measurable boost in visibility and relevance beyond what the acquirer's existing technologies experience.

These results again show that the citation gains following acquisition are not merely part of a broader firm-wide trend. Even when compared directly to the acquirer's own innovation output, acquired patents receive a larger citation boost.


Table A.3: Control group: Patents by acquirer

	(1)	(2)
	#Citations	#Citations
Post	0.92*** (0.10)	0.88*** (0.09)
Treated	0.31*** (0.05)	
Age ²		-0.02*** (0.00)
Age ³		0.00*** (0.00)
CPC \times Year FE	yes	yes
Patent ID FE	no	yes
Age FE	yes	no
Adj. R2	0.07	0.45
Avg. Outcome	0.98	0.98
N	2978341	2978341

Standard errors in parentheses

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

Figure A.4: Event study, Patents by acquirer



A.3 Additional Figures and Tables

This section provides supplementary figures and tables that complement and extend the main results presented in the paper.

Table A.4: Dealsize, patents and target size

	(1) Dealsize (M\$)
Target with Patents	468.99** (227.52)
# Employees	0.58*** (0.00)
Target with Patents \times # Employees	1.25*** (0.08)
Constant	246.92*** (76.43)
Observations	198

Robust standard errors in parentheses

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

Figure A.5: Patent family citations




Figure A.6: Years to next own acquisition within the same CPC group

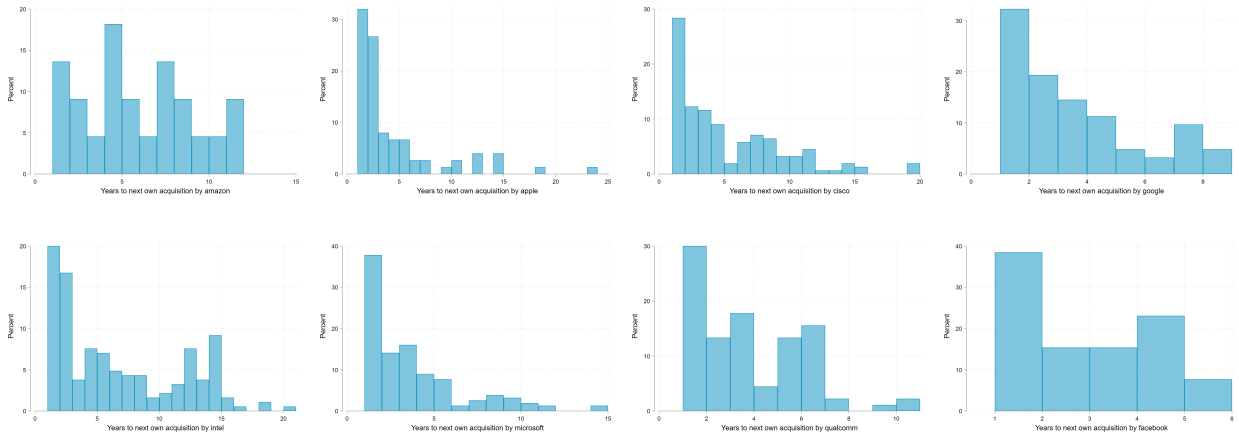


Figure A.7: Years to next acquisition, addressing selection

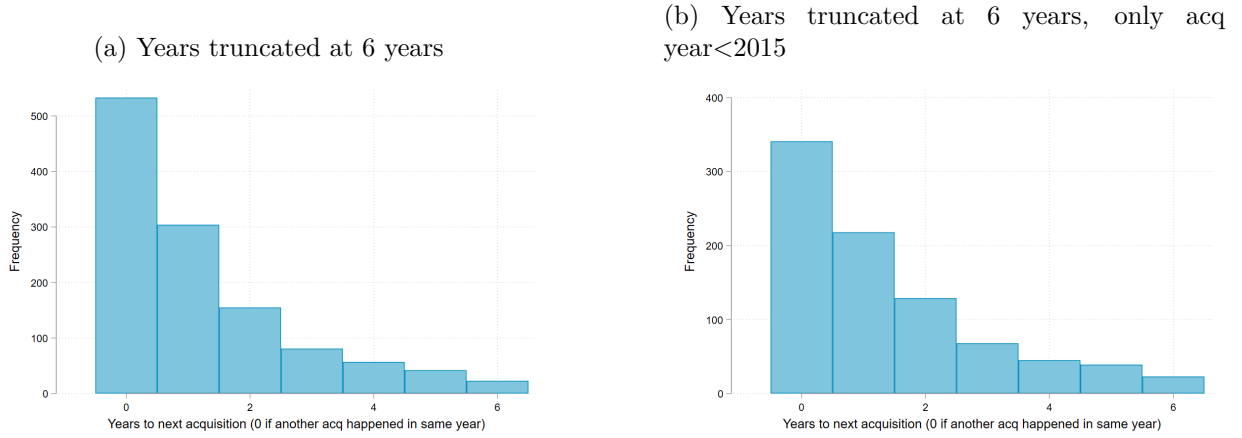


Figure A.8: First acquisition events within CPC group




Figure A.9: Matched event studies; All first acquisition events




Figure A.10: Retention rates by acquirer

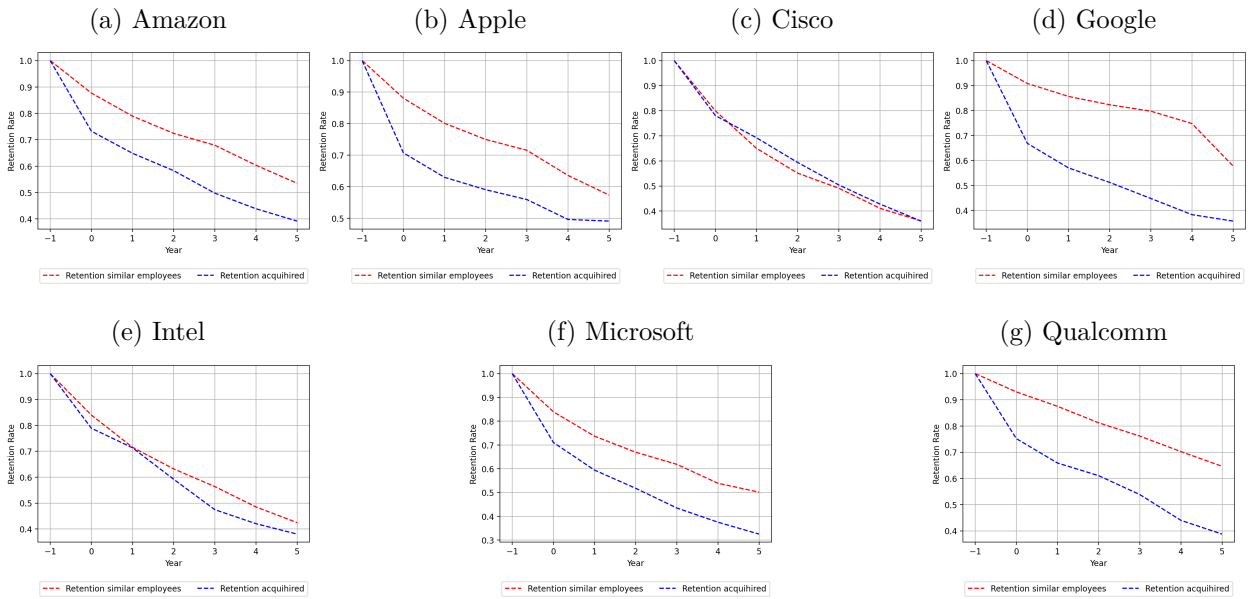


Table A.5: Innovation Heterogeneity by Number of Events

	(1) Unmatched	(2) Matched sample
Post 1st Event \times NoFollowOn	0.04 (0.09)	-0.15 (0.10)
Post 1st Event \times FollowOn	0.48*** (0.05)	0.35*** (0.06)
Post 2nd Event \times NoFollowOn	-0.12 (0.12)	-0.21 (0.13)
Post 2nd Event \times FollowOn	0.22*** (0.05)	0.17*** (0.05)
Post 3rd Event \times NoFollowOn	-0.06 (0.14)	-0.12 (0.14)
Post 3rd Event \times FollowOn	0.07 (0.06)	0.07 (0.07)
Post 4th Event \times NoFollowOn	-0.19 (0.18)	-0.16 (0.17)
Post 4th Event \times FollowOn	0.17** (0.08)	0.05 (0.11)
Constant	2.75*** (0.02)	4.62*** (0.04)
CPC FE	yes	yes
Year FE	yes	yes
Adj. R2	0.12	0.23
N	34236	14850

Robust standard errors in parentheses * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The table shows the result of a difference-in-difference regression specification where we estimate the marginal effect of a k th acquisition event on overall patenting activity depending on whether the affected CPC group experiences k (NoFollowOn) or more events(FollowOn) in total.

Table A.6: Control group: All patents in treated CPC groups - Grant year citations


	(1) #GYCitations	(4) #GYCitations	#GYCitations	#GYCitations
PostAcquisition	1.38*** (0.09)	1.48*** (0.09)	1.18*** (0.09)	1.30*** (0.09)
Acquired	0.13*** (0.03)		0.11*** (0.03)	
Post \times FirstNoFollowon			0.38 (0.44)	0.59 (0.47)
Post \times FirstFollowon			1.32** (0.53)	0.97* (0.55)
CPC \times Year FE	yes	yes	yes	yes
Patent ID FE	no	yes	no	yes
Age FE	yes	no	yes	no
Adj. R2	0.06	0.38	0.06	0.38
Avg. Outcome	0.72	0.72	0.72	0.72
N	44252928	44252928	44252928	44252928

Clustered standard errors in parentheses

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

Figure A.11: Effect on patent citations, below and above median retention

(a) Application year citations



(b) Grant year citations

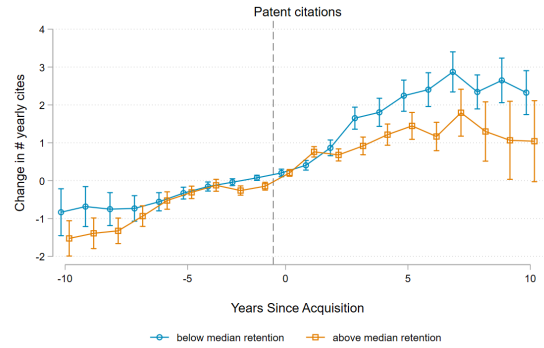


Table A.7: Citation effect by retention rate at the acquirer - Grant year citations

	(1)	(2)
	#GYCitations	#GYCitations
PostAcquisitions	1.21*** (0.09)	1.34*** (0.10)
Post×Retention	-1.44*** (0.42)	-0.96** (0.43)
Acquired	0.05* (0.03)	
Acquired×Retention	0.83*** (0.16)	
CPC × Year FE	yes	yes
Patent ID FE	no	yes
Age FE	yes	no
Adj. R2	0.06	0.38
Avg. Outcome	0.74	0.74
SD Retention	0.19	0.19
N	41878689	41878689

Clustered standard errors in parentheses

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