

# Being acquired and innovation : an empirical study

Dylan Alezra\*      Benoît Berquier†

December 23, 2024

## Abstract

This study examines the impact of acquisitions on innovation output in acquired firms, addressing concerns about potential negative effects on market competition and technological progress. Using a comprehensive dataset combining Crunchbase information and a custom-built patent database, we analyze innovation patterns before and after acquisitions. We establish key stylized facts: acquisitions have become the primary exit strategy for firms, surpassing IPOs; acquired firms tend to be technologically and industrially relevant to their acquirers; and there is often a decline in innovation within acquired firms post-acquisition although they innovate more on average in their early years.

Using propensity score matching and a difference-in-differences approach, our results show that acquisition leads to a substantial decrease in innovation for acquired firms. Our baseline Poisson regression reveals a 27.2% decrease in patent output following acquisition, which remains significant at 23.4% even after controlling for firm closures. This effect operates through both extensive and intensive margins: acquisitions increase the probability of completely ceasing innovation by 10.3 percentage points and significantly reduce the likelihood of high-volume patent production. While firms are also 8.6 percentage points more likely to close post-acquisition, our findings indicate that the decline in innovation extends beyond mere firm closures. These results suggest potential competitive concerns, although competition policy should consider both the *ex post* decline in innovation and the *ex ante* incentives acquisitions create for start-up innovation and market entry, rather than focusing solely on post-acquisition effects.

*Keywords:* innovation, killer acquisition, competition.

*JEL Code:* O30, G34, L40, O31, C23, L26

---

\*University of Paris-II Assas-Panthéon (CRED)

†University of Paris-I (CEPN)

# 1 Introduction

Since the 2010s, acquisitions have emerged as a seemingly safe route to innovation for incumbent companies across various markets, particularly in the technology sector. This strategic choice may help alleviate the innovator’s dilemma (Christensen, 1997). Indeed, firms may opt to purchase innovative start-ups or competitors to avoid the potential sunk costs associated with developing new innovations internally, which could potentially come at the expense of their current business model. This approach has led to a significant increase in the number of acquisitions. For instance, according to the Federal Trade Commission (FTC), the GAFAM companies (Google, Apple, Facebook, Amazon, and Microsoft) averaged between 70 and 100 acquisitions annually from 2010 to 2019 (FTC, 2021).

At the same time, the prospect of being acquired is often viewed as a success for new firms and may have served as an incentive for market entry and innovation (Eisfeld, 2024; Bisceglia et al., 2023; Cabral, 2021; Phillips and Zhdanov, 2013; Rasmusen, 1988). Consequently, this dynamic may have been favorable to innovation, at least in theory, by creating incentives for founders and helping incumbents stay ahead of cutting-edge technologies. However, this trend has also raised numerous concerns among academics and regulators. Cunningham et al. (2021) highlighted that a portion of these acquisitions are actually intended to eliminate potentially disruptive technologies and prevent the emergence of competitors, and coined the term “killer acquisition”<sup>1</sup>. On the political side, competition authorities worldwide have become acutely aware of transactions made by incumbent firms that fall below notification thresholds. A landmark case is the European Commission’s prohibition of Illumina’s acquisition of GRAIL in 2022, which was labeled for the first time as a “killer acquisition”. Despite these concerns, empirical evidence on the prevalence of such acquisitions remains very limited (Ivaldi et al., 2023; Gautier and Lamesch, 2021; Sokol, 2020). Therefore, should we be worried about this trend in acquisitions on innovation?

To address these concerns, our research question specifically is: does acquisition lead to a decrease in innovation for acquired firms?<sup>2</sup> Our main goal is to evaluate whether there is a decrease in the ability or incentive to innovate for these firms. Compared to previous studies, we provide a holistic and generalizable result, focusing not on a particular sector but on the impact of acquisitions on innovation across acquired firms regardless of

---

<sup>1</sup>This concept is not entirely new, as Ivaldi et al. (2023) point out earlier definitions of killer acquisitions, such as that by Carlin et al. (2007), who described them as “transactions that lead not just to value dissipation for the acquiring parties, but that result in such a profoundly negative outcome that the fact of the consummation of the transaction in fact results in the onset of financial distress and potential liquidation for the newly-enlarged firm”. Separately, historical examples of such practices include cases like Standard Oil buying out tramway companies to limit their development and favor the growth of car usage.

<sup>2</sup>While we acknowledge that the overall impact of acquisitions on innovation should be evaluated at a sectoral level, as a potential decrease in innovation by acquired firms might be offset by increased innovation elsewhere in the sector or by the acquirer, it is crucial to first assess the direct impact on acquired firms.

their industry. Indeed, while much research has focused on acquisitions in software and pharmaceutical industries due to their frequency and innovation intensity, we believe it is important to understand the general relationship between acquisitions and innovation at a broader level, and avoid what has been criticized as a discriminatory focus on a few industries (Yun, 2022)<sup>3</sup>.

In order to answer this question, we have constructed a comprehensive and unique dataset that overcomes many traditional limitations for observability. Our data collection process draws from two main sources. First, we utilize all available data from Crunchbase, including firm characteristics, textual descriptions of companies, complete lists of acquisitions and IPOs, and information on funding rounds. This data encompasses more than 3 million firms, approximately 140,000 acquisitions, 600,000 funding rounds, and 40,000 IPOs. Secondly, we have built our own patent database as a free alternative to PatStat. This approach allows us to expand beyond the traditionally used USPTO dataset and include patent data from all available jurisdictions since the 1800s.

Our first contribution is related to industry definition using textual data. Because the Crunchbase industry classification can be limited, not offering a simple way to define sectors and potential markets, we have developed an alternative approach. Using text embeddings and cosine similarity measures, we are able to match firms to their NAICS sectors based on their descriptions. This method not only overcomes the limitations of the Crunchbase classification but also enables us to analyze similarities between firms. Indeed, we leverage this approach to identify potential sectoral overlaps between acquirers and acquired firms, defined by their semantic similarity (or an overlap in their Crunchbase industry classification). Additionally, we utilize our patent dataset to assess technological similarity. We define technological similarity as the overlap between International Patent Classifications, at the most precise level, in both the acquired and acquirer’s portfolios in the pre-acquisition period.

We document the trends in acquisitions since the 1990s, on a large scale, extending previous research which had focused on specific industries. Acquisitions as an exit strategy have tripled in 20 years. Acquired firms are often close to their acquirer, both sectorally and technologically. However, as often shown, being acquired is often associated with closure, and less innovation in the long run, especially compared to non-acquired firms.

In this paper, we provide new evidence of a substantial negative impact of acquisitions on innovation for acquired firms. To our knowledge, we are the first to provide evidence of the effect of acquisition on acquired firms, regardless of firm type or sector, and to provide a comprehensive study of the effect. To examine the effects of acquisitions on innovation, we employ both baseline Poisson regression and a difference-in-differences (DiD) analysis

---

<sup>3</sup>Given the nature of our data, however, it is biased toward technological firms.

with a matched control group. Our baseline Poisson regression reveals that being acquired is associated with a 27.2% decrease in patent output, which remains significant and of a similar magnitude after controlling for firm closures. Our analysis shows that this decline operates through mostly the extensive margin but also to the intensive margin to some degree. Acquisitions increase the probability of completely ceasing innovation by 10.3 percentage points and the likelihood of closure by 8.6 percentage points. Moreover, acquired firms show significantly reduced probabilities of producing higher numbers of patents, with decreases of up to 3.17 percentage points in the likelihood of producing 3-5 patents and 3.15 percentage points for producing more than 10 patents. These findings indicate that acquisitions not only increase the likelihood of firms ceasing innovation entirely but also shift surviving firms toward lower levels of innovative output.

**Literature Review** Our main contribution to the literature relates to the relationship between market structures and innovation. [Schumpeter \(1942\)](#) posits that entrepreneurs innovate to benefit from temporary monopoly profits as a way to escape competition. Consequently, there is a creative destruction happening as previous incumbents are displaced by new ones through innovation. His full thesis, however, suggests that in the long run, firms' bureaucratization and leading incumbents' attempts to automatize technological progress hamper the creative destruction process through acquisition and ultimately reduce incentives for market entry and entrepreneurship.<sup>4</sup>

Therefore, it may explain through a simple rationale the role of acquisition in the current landscape. Incumbents, in technological or pharmaceutical markets for instance, may have used acquisitions as an R&D tool ([Higgins and Rodriguez, 2006](#)), and may have created "kill zones" ([Kamepalli et al., 2021](#)). Yet this prevalence of monopolies may be detrimental to innovation overall ([Arrow, 1962](#)). Theoretically, acquisitions can be pro-competitive when the acquirer helps the target's innovative efforts, or anti-competitive if the acquired firm could have pursued innovation independently ([Affeldt and Kesler, 2021a](#); [Brutti and Rojas, 2022](#); [Hollenbeck, 2020](#); [Jullien and Lefouili, 2018](#); [Federico et al., 2018](#); [Mermelstein et al., 2014](#)). However, several studies have shown the negative impact of M&A on innovation ([Poege, 2022](#); [Fons-Rosen et al., 2021](#); [Watzinger et al., 2020](#); [Haucap et al., 2019](#); [Seru, 2014](#)). In particular, a number of papers have investigated the effect of large incumbent acquisitions, especially by GAFAM, on innovation and outcomes correlated with innovation,

---

<sup>4</sup>"Since capitalist enterprise, by its very achievements, tends to automatize progress, we conclude that it tends to make itself superfluous—to break to pieces under the pressure of its own success. The perfectly bureaucratized giant industrial unit not only ousts the small or medium-sized firm and 'expropriates' its owners, but in the end it also ousts the entrepreneur and expropriates the bourgeoisie as a class which in the process stands to lose not only its income but also what is infinitely more important, its function. The true pacemakers of socialism were not the intellectuals or agitators who preached it but the Vanderbilts, Carnegies and Rockefellers." ([Schumpeter, 1942](#), p. 134)

notably patents. While there is not necessarily a consensus, the literature does not suggest an overall positive effect in the long run of acquisitions on innovation, but may depend on the type of acquirers or the sector (de Bary and Gautier, 2024; Prado and Bauer, 2022; Gugler et al., 2023).

Our article adds to this empirical literature, allowing us to show a clear negative effect on patent filing, even after controlling for firm closure. The specificity of our paper is the scale at which we observe this effect. Indeed, we detail the overall effect for all firms acquired. Moreover, we believe that our methodology is helpful when it comes to finding a direct counterfactual for an acquired firm, especially using panel data. Consequently, this article also relates to the rich theoretical literature on the relationship between acquisitions and innovation (Letina et al., 2021; Hollenbeck, 2020; Jullien and Lefouili, 2018; Federico et al., 2018; Mermelstein et al., 2014), providing empirical evidence of some of the predicted effects.

However, as evidenced by the literature, there is a clear *ex ante* incentive for market entry and therefore potentially beneficial to innovation overall (Eisfeld, 2024; Warg, 2021; Wang, 2018)<sup>5</sup>. Start-ups, for instance, may want to cater to incumbent innovation in order to be bought later on. This paper can provide evidence, at the descriptive level, for the *ex ante* incentive to innovate in order to be acquired. Indeed, statistically, future acquired firms innovate more in their early years, both quantitatively and qualitatively, compared to non-acquired firms. The difference in trends in the long run observed between the two groups may be explained by this pre-acquisition incentive.

**Outline** The paper is structured as follows: Section 2 details the data construction process. Section 3 presents stylized facts on acquisition trends since the 1990s. Section 4 details our empirical strategy. Section 5 presents our results. Section 6 discusses the results. Section 7 concludes the paper.

## 2 Data

### 2.1 Using Crunchbase data

Crunchbase<sup>6</sup> is a comprehensive data repository providing financial and operational information on over 3 million companies worldwide, encompassing both public and private entities. For approximately one-third of these entities, Crunchbase offers a detailed description, while the remainder are accompanied by short descriptions. Whenever possible, we use the long description as it provides more information on the nature of the business.

---

<sup>5</sup>Unfortunately, we haven't found an empirical study that compares the magnitude of the *ex ante* effect and the *ex post* effect on innovation

<sup>6</sup>[www.crunchbase.com](http://www.crunchbase.com)

The database captures a wide range of corporate events and characteristics, including founding dates, funding rounds, acquisitions, investments, initial public offerings (IPOs), and company closures. Crunchbase also employs its own proprietary industry classification system, which serves to further characterize the activity of a firm, although it has its limitations.

Using this information, we construct two datasets. First, we build a large cross-sectional database by utilizing Crunchbase’s unique URL identifiers for each organization, matching our different cross-sections on acquisitions, funding rounds, and IPOs. Second, we build a longitudinal panel dataset by exploiting founding dates and, where applicable, closure dates. We ensure to keep only firms where relevant information is not missing, such as the founding date and the number of employees. This reduces the dataset to 2 447 991 firms, which remains an incredible amount of data for our purposes. When the closing date isn’t available for an acquired firm, we assume it was closed at the time of the acquisition, unless patent data proves otherwise.

Table 1: Statistical Summary, per year, between 1990 and 2023

Category	Min	Max	Median	Observations
Acquisitions	19	14,094	1,912	149,342
IPOs	188	2,364	1,087	45,388
Funding Rounds	48	63,524	18,085	615,111
Patent Filing Firms	15,225	69,483	41,431	161,011

Table 1 presents a summary of the data collected. While Crunchbase’s coverage of venture capital activities may not be as exhaustive as specialized platforms like Pitchbook, it offers a significant advantage in terms of accessibility. This characteristic has contributed to its widespread use in empirical finance and competition research (Eisfeld, 2024; Gugler et al., 2023; Denes et al., 2022).

### 2.1.1 Defining relevant variables

**NAICS classification using Mistral AI embeddings** To define markets, identify potential competitors more precisely, and improve industry classification, we leverage Mistral AI embeddings. We provide a detailed rationale for this choice in the appendix. For each company  $i$ , we create an embedding  $e_i$  based on a concatenation of the company name, Crunchbase industries, and full description:

$$e_i = f_{\text{Mistral}}(d_i)$$

where  $f_{\text{Mistral}}$  is the Mistral AI embedding function that maps the input to a high-dimensional vector  $e_i \in \mathbb{R}^n$ .

We use the embeddings to assign the two most probable 6-digit NAICS codes to each company by comparing company embeddings to pre-computed NAICS category embeddings. This method offers several advantages over traditional industry classification systems, including Crunchbase’s native industry tags. Indeed, the Crunchbase classification is limited in that companies operating in similar domains (e.g., veterinary-related activities) may not be classified in the same industry or considered proximate. Our approach addresses this limitation by using semantic similarity.

For large incumbents who operates on many markets (such as Google), it does not necessarily provide the best categorization. However, we build a typology of incumbents, which we believe to be a solid complementary approach to classify firms. We provide a distribution of acquisitions by 2-digit NAICS codes to exemplify our approach in Figure 1. As expected, most acquisitions tend to be in industrial and scientific sectors.

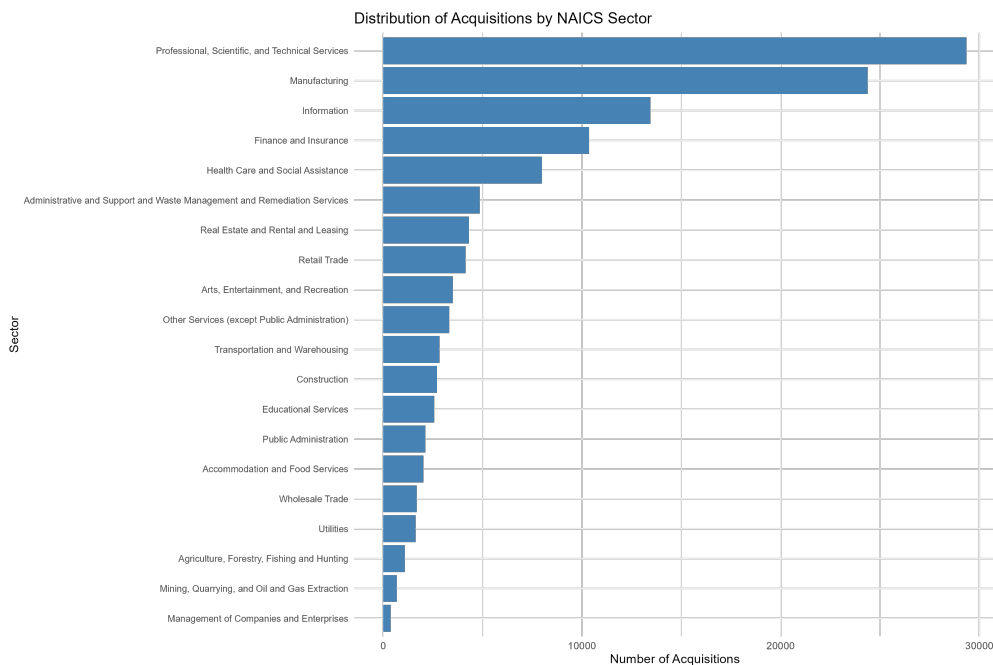


Figure 1: This figure represent the number of acquisition by 2-digit NAICS codes

**Overlapping** Measuring overlapping between acquired firms and acquirers is a common practice to evaluate the similarity between them, which may be done at the project or product level (Cunningham et al., 2021; Einfeld, 2024). While our data doesn’t allow us to go at such a precise, we define a sectoral overlapping measure. It is a dummy variable that equals 1 if an acquiree has in common at least one industry (either defined by Crunchbase, or by NAICS), or a semantical similarity higher than 0.7 using their embedded descriptions.

We also isolate the patent class measure to define a technological overlap<sup>7</sup>. This restriction allows us to delve further into the analysis, as it is more precise measure of similarity at the technological, even if it is limited by the number of observations.

**Acquirer Types** We classify acquirers into four main categories, following partly [Eisfeld \(2024\)](#):

- *GAFAM*: Google, Amazon, Facebook, Apple, and Microsoft.
- *Old Incumbents*: Firms with over 10,000 employees or that have gone public, established before 1995.
- *New Incumbents*: Similar to Old Incumbents, but established after 1995.
- *Others*: Firms not falling into the above categories, including pre-exit firms (those that haven't exited or reached a significant size) and hard-to-classify firms that also make acquisitions.

For reference in our descriptive statistics, we also identify two additional groups, that have been often ignored in the recent literature:

- *NATU*: Netflix, Airbnb, Tesla, and Uber.
- *BATX*: Baidu, Alibaba, Tencent, Xiaomi, and other significant Chinese firms such as Didi Chuxing, Huawei, ByteDance, and JD.com.

## 2.2 Patent data

### 2.2.1 Building a patent database from scratch

To our knowledge, no comprehensive, publicly available, and ready-to-use patent database exists. Freely accessible options are either unstructured (lens.org, Google Patents) or incomplete (U.S. Patent Office, restricted to U.S. jurisdiction). Paid alternatives, while comprehensive, are prohibitively expensive (e.g., PatStat) and often difficult to use.

Consequently, we opted to construct our own patent database by web scraping and downloading patents from Google Patents and lens.org, creating an open, structured patent database. We detail the database construction, cleaning process, entity linking, and named entity recognition in the appendix. We have named this database OpenPat to facilitate researches on innovation. OpenPat contains over 130 million observations, including patent

---

<sup>7</sup>We use the most precise level of the IPC class



names, abstracts, jurisdictions, families, International Patent Classifications (IPC), publication dates, and patent statuses. It encompasses all publicly available patent data from 1800 to 2024, across all jurisdictions, in a structured and user-friendly format.

### **2.2.2 Matching Patent data to Crunchbase**

A significant empirical challenge lies in matching Crunchbase data to our patent data. We developed a specialized algorithm combining exact matching, fuzzy matching, and embeddings to address this issue. The detailed process is outlined in the appendix. We successfully matched approximately 161,000 entities, though future research may yield additional matches. However, there exists a trade-off between match accuracy and the number of potential matches.

### **2.2.3 Patent variables**

To measure innovative activity, we utilize three variables derived from our patent data. The first variable, patent count, represents the raw number of patents. Since the same innovation can be protected by different patents across jurisdictions, it may also represent a firm’s willingness to protect an innovation. Our second variable weights patents by their family size. This allows us to capture the actual number of innovations produced by a firm and avoid overestimation due to firms seeking broad geographical protection (Squicciarini et al., 2013). This is the variable we use for our regression analysis.

Across all these measures, we consistently use the patent application date rather than the publication date, as this more accurately reflects the timing of a firm’s innovative activity, especially given the often significant delay between application and publication. For some of our stylized facts and some regressions, we use the log number of weighted patents due to extreme values, which appears to be a common practice in the literature (Gugler et al., 2023; de Barys and Gautier, 2024). However, later on, we explain the interpretability problem with the use of log-transformed variables.

## **3 Stylized facts**

### **3.1 Acquisitions have become the primary exit strategy for firms in recent years**

The number of acquisitions has increased significantly since the 2000s, with a striking surge of 156% between 2010 and 2023. Meanwhile, the number of IPOs has decreased by 64.8% over the same period (Figure 2).

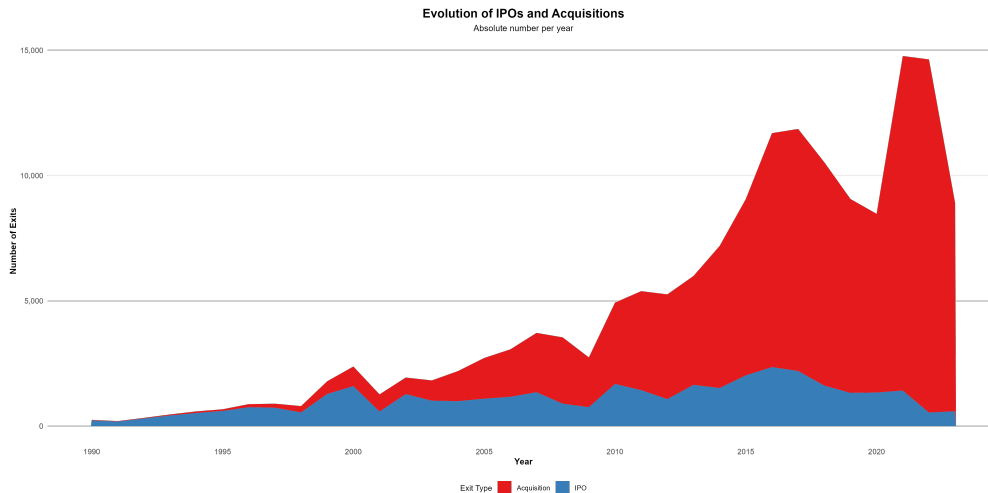


Figure 2: This figure shows the annual number of exit events through IPOs (in blue) and acquisitions (in red) for firms between 1990 and 2023, using all data available in Crunchbase.

This trend has led to a latent debate in the literature. On the one hand, [Cunningham et al. \(2021\)](#) have emphasized the predatory behavior of incumbents, which may have hampered firm growth, using the pharmaceutical industry as an example. However, this is not a sufficient explanation given the low estimated percentage of so-called “killer acquisitions”. On the other hand, [Lemley and McCreary \(2019\)](#) suggests it reflects an evolution in exit strategies rather than predatory behavior.

For venture capitalists, exits typically occur through IPOs or acquisitions for the firms they fund. In a landscape where incumbents seek acquisitions to enhance their products and have substantial cash reserves, while entrepreneurs may also want to be rewarded, this coincidence will lead to premature acquisition exits instead of going through the IPO, which can be viewed as a firm’s ultimate success in terms of growth. They also hypothesize that given the uncertainty of the firm’s success or failure, and as it takes more time to get to an IPO, VCs may pressure entrepreneurs to exit more rapidly through an acquisition.

Our data corroborate Lemley’s observations, showing a preference for acquisitions over IPOs among VC-funded companies (Figure 3). This trend does not seem to be influenced solely by venture capitalists, while being more pronounced among VC-backed firms. This trend is actually a global phenomenon since the 2000s. It aligns with the literature arguing that the prospect of being acquired may motivate an entry on a market ([Warg, 2021](#)).

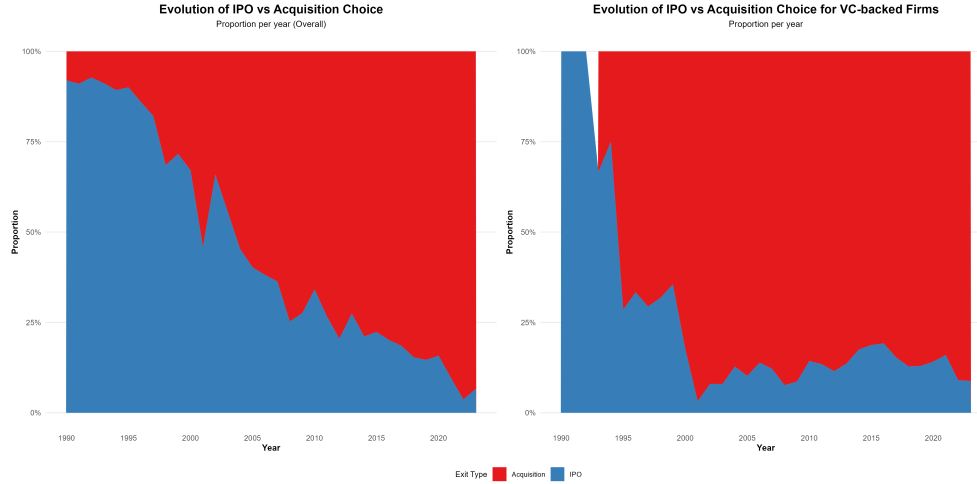


Figure 3: This figure presents the relative proportion of exits through IPOs versus acquisitions from 1990 to 2023, using Crunchbase data. The left panel shows the evolution for all firms, while the right panel focuses specifically on VC-backed firms. Both graphs display the ratio of acquisitions to total exits (acquisitions + IPOs) over time. Note that data before 1995 should be interpreted with caution due to potential coverage limitations in the Crunchbase database for this early period.

Moreover, the time to exit does not seem to correlate with this trend. Indeed, the median time to exit has remained relatively stable across funding modes (Figure 4). Actually, it is worth noting that not only is the median time to exit for VC-backed firms significantly lower than for non-VC-backed firms, but also that within VC-backed firms, there is not such a striking difference in the median time.

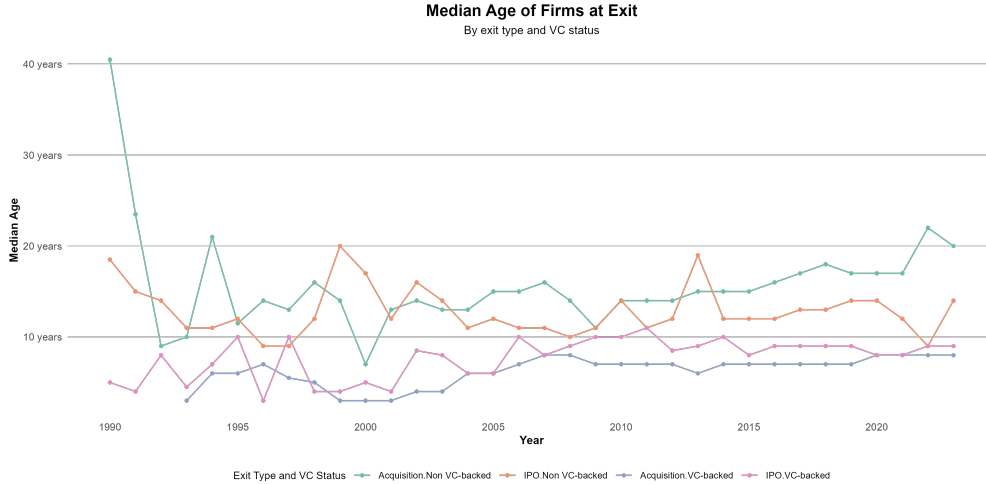


Figure 4: This figure displays the median age of firms at exit from 1990 to 2023, differentiated by exit type (IPO or acquisition) and VC funding status. The graph shows four categories: IPOs of non-VC-backed firms (green), acquisitions of non-VC-backed firms (orange), IPOs of VC-backed firms (purple), and acquisitions of VC-backed firms (pink). The median age is calculated as the difference between the exit year and the founding year. Note that data before 1995 should be interpreted with caution due to potential coverage limitations in the Crunchbase database for this early period, particularly for non-VC-backed firms.

In order to confirm those intuitions, at least descriptively, we use a multinomial regression to model the exit choice between IPO and Acquisition (using IPO as referential). Table 2 presents our results. The negative coefficient on VC-backing across all models indicates that VC-backed firms are generally less likely to exit via IPO compared to acquisition. However, this effect is moderated by several factors. Higher VC intensity (Model 2) slightly mitigates the negative impact of VC-backing on IPO likelihood. Interestingly, the post-2000 era (Model 3) sees a significant increase in the probability of IPO for VC-backed firms, despite the overall trend towards acquisitions<sup>8</sup>. Finally, as VC-backed firms age (Model 4), their probability of choosing an IPO exit marginally increases, possibly reflecting a maturation effect. It still does not provide evidence for the thesis that VC-backed firms are pressured into selling to an incumbent.

<sup>8</sup>The effect of the period however is likely due to the higher data availability on VC-backed firms in our sample, even though there is a so-called effect of the period

Table 2

	<i>Dependent variable:</i>			
	Exit Type (IPO vs Acquisition)			
	Base Model	VC Intensity Interaction	Post-2000 Interaction	Age Interaction
	(1)	(2)	(3)	(4)
VC-backed	-2.326*** (0.033)	-1.560*** (0.012)	-3.543*** (0.017)	-2.516*** (0.040)
VC Intensity	0.300*** (0.007)	0.713*** (0.015)	0.298*** (0.007)	0.301*** (0.007)
Post-2000	-0.976*** (0.001)	-1.014*** (0.001)	-1.511*** (0.0004)	-1.002*** (0.001)
VC-backed * VC Intensity		-0.456*** (0.012)		
VC-backed * Post-2000			1.234*** (0.017)	
VC-backed * Age at Exit				0.017*** (0.002)
Controls	Yes	Yes	Yes	Yes
Observations	166,690	166,690	166,690	166,690

*Note:*

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01

### 3.2 Acquired companies tend to be technologically and industrially relevant to their acquirers

While the literature on acquisitions has often focused on start-ups or specific industries, our study aims to provide a more comprehensive analysis of the acquisition phenomenon.<sup>9</sup>

<sup>9</sup>Part of the motivation of this article is associated with the literature on killer acquisitions. It is often defined as an incumbent buying promising start-ups in order to prevent them from becoming a serious competitor within a given market. However, a killer acquisition, at least theoretically, may be defined more accurately as buying a potential or actual competitor with some degree of overlapping, regardless of the acquired firm size or funding type.

Table 3

Characteristic	Overall	Missing (%)
<i>Acquired firms characteristics</i>		
Size (median)	11-50	29.1
Relative age (median)	14	13.8
Patent owners (%)	5	0
Patent applicants (%)	8.6	0
VC-backed (%)	76.7	77.4
Number of rounds (median)	2	77.4
Created after 2000 (%)	50.8	13.8
<i>Compared to acquirer</i>		
Sectoral similarity (%)	64.1	4.2
Acquiree is younger (%)	61.7	26.7
Acquiree is smaller (%)	75.7	36.6
Patent overlap among owners (%)	31.8	0
Total observations	149,301	

*Note:* Patent overlapping was computed on the basis of all acquirers and acquiree's that possess patent.

Table 3 presents a detailed overview of acquired firms' characteristics. As anticipated, acquired firms tend to be relatively young and small. The median acquired firm has only 11 to 50 employees and is 14 years old, with 50.8% of targeted firms created after 2000. This trend towards acquiring younger companies is further emphasized by the data on VC-backed firms. Among VC-funded acquired firms, the median number of funding rounds is only 2, suggesting that many acquired firms are still in their early stages of development.

Acquired firms tend to be more innovative compared to the average. Interestingly, there's a distinction between patent owners and patent applicants among acquired firms. While 5% of acquired firms are patent owners, a larger proportion (8.6%) are patent applicants. This suggests that some acquired firms are actively pursuing innovation even if they haven't yet secured patents. Both these figures likely exceed the average innovation rates in the complete database (around 5% of patent applicants in the total database), highlighting the innovative nature of acquisition targets.

The most striking feature of these acquisitions is their technological and industrial relevance. Indeed, among transactions involving patent owners, the degree of technological overlapping is 37.2%. While lower than previously reported, this still indicates a substantial

alignment in innovation areas between acquirers and acquirees. This finding can be extended to industrial relevance, with a high degree of sectoral similarity (64.1%).

Acquirers tend to buy firms that innovate in similar areas as theirs. The data also shows a tendency to acquire younger (61.7%) and smaller firms (75.7%), suggesting that acquired firms are more often than not a means to refine the competitive edge of acquirers.

Table 4

Characteristic	GAFAM	Old Incumbent	New Incumbent	Others
<i>Acquired firms characteristics</i>				
Size (median)	11-50	11-50	11-50	11-50
Relative age (median)	5	10	17	14
Patent owners (%)	6	5.1	7.3	3.9
Patent applicants (%)	18.4	9.5	12.3	6.5
VC-backed (%)	94.9	82.3	75.2	75.4
Number of rounds (median)	2	2	2	2
Created after 2000 (%)	76.7	62.5	35.2	55.9
<i>Compared to acquirer</i>				
Sectoral similarity (%)	47.8	69.2	66.3	61.9
Acquiree is younger (%)	94.3	50.6	87.1	51.4
Acquiree is smaller (%)	95.1	83.6	89.4	66.9
Patent overlap among owners (%)	62.6	29.6	34.8	19.7
Total observations	761	18,542	42,650	87,348

*Note:* Patent overlapping was computed on the basis of all acquirers and acquiree's that possess patent.

Table 4 shows distinct patterns across acquirer types. New incumbents tend to purchase older firms (median age 17) compared to old incumbents (median age 10), while GAFAM targets are typically the youngest (median age 5). A common trend emerges across all acquirers: targets are generally younger and smaller, with 87.1% to 95.1% of acquisitions involving younger firms and 66.9% to 95.1% involving smaller entities. Patent activity is particularly high among acquired firms, with GAFAM targets showing the highest rates of patent ownership (6%) and applications (18.4%). Notably, GAFAM acquisitions stand out with the highest patent overlap (70.5%) despite lower sectoral similarity (47.8%) compared to other acquirers (61.9% to 69.2%). It may give credit to the thesis supporting that GAFAM are buying potential rivals, but it is a question of interpretation, as it might as well be a way to enhance their current technologies. Also, on average, per firm, GAFAM are known to buy out way more firms than other incumbents. The prevalence of VC-backing in acquired

firms (75.2% to 94.9%) further underscores the focus on innovative, high-potential targets across the acquisition landscape. As shown in the appendix, BATX and NATU show similar behaviors to GAFAM, except they acquire fewer firms (Table 11).

### 3.3 There is often a decline in innovation within acquired firms following their acquisition

How does acquisition affect a firm’s innovative effort? The two relevant outcomes to investigate within our data are firm closure and innovative behavior (measured by patent appliance, patenting rate, and patent influence).

Table 5

(a) Overall Closure Rates

	Overall	Closed at Exit
Closed (%)	20	98.9
Total Observations	149,301	147,672

(b) Closure Rates and Observations by Sectoral and Patent Overlap

	Sectoral		Patent	
	No Overlap	Overlap	No Overlap	Overlap
Closed (%)	19	21.4	26.3	41.9
Total Observations	51,398	91,697	1,519	148

(c) Closure Rates by Acquirer Type

	GAFAM	New Incumbent	Old Incumbent	Others
Closed Rate (%)	47.3	23.5	20.8	18.6
Total Observations	761	18,542	42,650	87,348

Panel (a) of Table 5 illustrates closure rates overall and by categories. Approximately 20% of acquired firms cease operations, predominantly within the first year post-acquisition. This phenomenon may lead to an immediate drop in innovation, though this interpretation is naive and is discussed in the following subsection. Panel (b) reveals that firms with technological overlap with their acquirers are significantly more likely to close (36.8% vs. 26% without overlap), potentially diminishing innovation in these specific technological areas. Finally,



as shown in Panel (c), GAFAM close almost half their acquired firms, which is twice the average closure rate for all other types.

Table 6

Metric	Overall	Type of Firms			
	All Firms	GAFAM	New Incumbent	Old Incumbent	Others
<b>Patent Appliance</b>					
<i>Before</i>	8.57	18.40	9.47	12.31	6.47
<i>After</i>	4.85	11.83	5.40	8.12	3.07
<b>Patent Count Rate</b>					
<i>Before</i>	1.76	1.53	1.30	2.80	1.34
<i>After</i>	1.21	1.37	0.80	2.15	0.83
<b>Patent Weighted Rate</b>					
<i>Before</i>	0.49	0.30	0.32	0.75	0.40
<i>After</i>	0.26	0.24	0.14	0.47	0.18
<b>Patent Influence</b>					
<i>Before</i>	1.54	1.42	0.94	2.41	1.23
<i>After</i>	0.22	0.28	0.13	0.45	0.12
<b>Total Observations</b>	149,301	761	18,542	42,650	87,348

Regarding patenting outcomes, we observe an overall decrease in innovation<sup>10</sup> in Table 6. Across all acquired firms, patent appliance rates drop from 8.57 to 4.85 post-acquisition. The patent rate also decreases from 0.24 to 0.13, suggesting a general decline in patenting activity. Patent influence shows a substantial decrease from 0.77 to 0.11, indicating a reduction in the impact of patents filed post-acquisition. The trend appears similar for patent protection, albeit a bit less pronounced.

Notably, this trend varies across acquirer types. GAFAM acquisitions show the highest pre-acquisition patent appliance (18.40) but also the largest absolute decrease (to 9.47). New incumbents and old incumbents follow a similar pattern of decline. Interestingly, while patent rates generally decrease, GAFAM acquisitions show a slight increase from 0.15 to 0.16, suggesting a possible intensification of patenting efforts in these cases.

<sup>10</sup>This assessment assumes patents as a measure of innovation, which is open to debate.

### 3.4 In the long-run, non-acquired firms tend to innovate more, but there might be an *ex ante* incentive for acquired firms

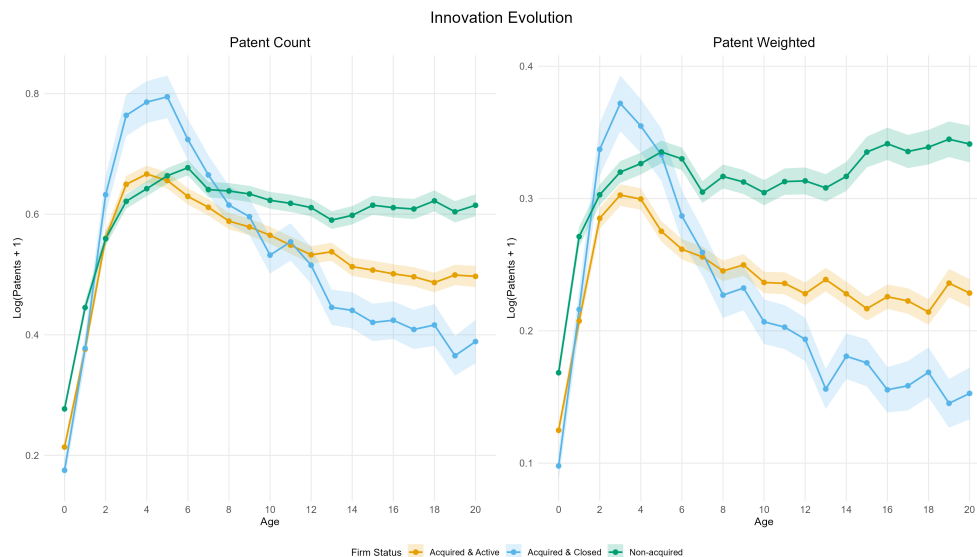


Figure 5: This figure shows innovation patterns over time for three types of firms using a balanced panel: non-acquired firms (N=5,967), acquired firms that remain active (N=5,045), and acquired firms that were subsequently closed (N=916). The left panel shows raw patent counts while the right panel displays patent-weighted measures adjusted for patent family size. Data is shown for the first 20 years of firm life and is log-transformed ( $\log(\text{Patents} + 1)$ ) to account for skewness.

Figure 5 shows the divergence in innovation between acquired and non-acquired firms over time. Initially, from founding to around the 7th-10th year, acquired firms exhibit higher innovation rates, particularly those that will eventually be closed, reaching their innovation peak early at age 3. During this early phase, acquired firms that will be closed show the highest patent counts, compared to both non-acquired and acquired but remaining active firms. This difference may reflect an *ex ante* incentive to innovate more to differentiate themselves from other firms, although it is not verifiable descriptively. A pivotal shift occurs around the median acquisition age, after which non-acquired firms demonstrate superior long-term innovation trajectories over the 20-year span. This trend suggests that firms achieving independent growth, such as those reaching IPO status like major tech incumbents, not only sustain but intensify their innovative efforts over time, as evidenced by their later peak age (17 years) compared to acquired firms that remain active or are eventually closed.

This last stylized fact appears to be quite significant in the context of our literature. While it could be said that acquired firms innovate less over time, it should be acknowledged that they contribute more to innovation, and possibly overall, than non-acquired firms in the

very short term. We investigate this fact further in Figure 6, which reveals a pronounced drop in innovation around the acquisition event. There is a clear temporal pattern around the acquisition suggesting that the decrease in innovation is causally linked with the acquisition process itself, rather than being driven by other confounding factors.

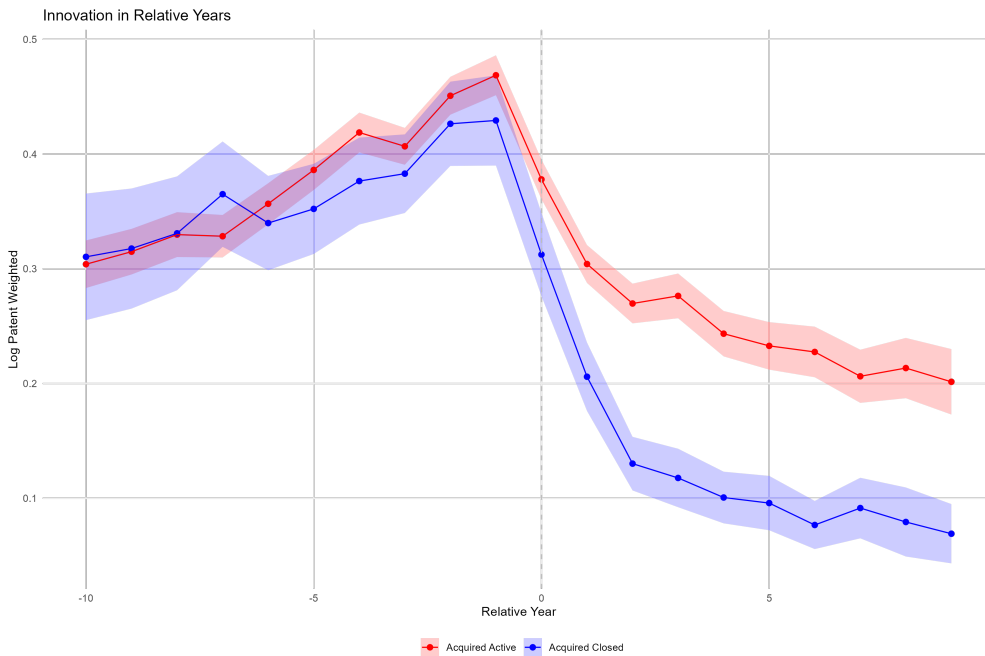


Figure 6: This figure plots innovation patterns around acquisition events (marked at relative year 0) for a balanced panel of acquired firms, comparing firms that remain active (N=5,344) versus those that were subsequently closed (N=950). Innovation is measured using weighted patent counts in logarithmic form. The sample spans from 10 years before to 9 years after acquisition.

## 4 Empirical strategy

### 4.1 Building a counterfactual

Without a counterfactual, and despite our descriptive evidence, it remains unclear whether the decline in innovative behavior is specific to acquired firms in the post-acquisition period or a common phenomenon among patent issuers. For instance, and quite counterintuitively, firms experiencing a downward trend in their innovation effort may be more likely to be acquired, in a somewhat similar fashion to what is observed by [Blonigen et al. \(2014\)](#) for productivity.

To address these potential selection biases and establish a reference point for comparing the drop in patent issuance, it is common practice to use a propensity score matching method

to find a control group of similar but non-acquired firms. Using our panel data for propensity score matching is not a conventional approach so we developed an ad hoc methodology to meet our matching requirements. We first conducted an initial matching to identify, for each treated firm (those with patents that were acquired), the 20 non-acquired patent-holding firms that were founded in the same year (or a year very closed) and operated in the same 3-digit NAICS sector. Then, for each matched firm, we define a fictive treatment date for non-treated firms by using the age of the treated firm at the moment of the acquisition in their matched group (leveraging the matching based on the founded year in the first stage). In the second stage of our matching procedure, we implement a more refined matching strategy focusing on pre-treatment patent trajectories. Specifically, for each treated firm, we select the most similar control firm within its first-stage matched group based on patent production patterns in the three years preceding acquisition ( $t - 3$  to  $t - 1$ ). The matching quality is assessed using standardized differences (t-statistics), with a conservative threshold of 0.05 to ensure strong pre-trend similarity. This two-stage approach allows us to construct a control group that closely mirrors both the sectoral characteristics and innovative dynamics of acquired firms prior to acquisition.

Our final matched panel is balanced around the acquisition event, spanning from three years before ( $t - 3$ ) to four years after ( $t + 4$ ) the acquisition (including the acquisition year). We obtain 4244 controls for our 4244 treated firms. Our matching rate is around 35%. There was a trade-off between the quality of the matching in the pre acquisition period and the matching rate. However, we wanted to ensure to most robust control group in order to respect the Parallel Trend Assumption<sup>11</sup>

We want to emphasize that it ultimately the quality of the matching relies on two crucial assumptions. First, firms that are founded the same year, operate in the same 3-digits NAICS sector and had the same pre-trends in term of real innovation defined as patent weighted variable are likely to be similar. Second, unobservables factors are correlated with the trend in patent production. Those are strong assumptions but allows to find what we believe to be a solid control group.

## 4.2 Estimation framework

Traditional approaches to analyzing patent data often employ log transformations to handle their skewed distribution. However, as [Chen and Roth \(2024\)](#) demonstrate, this practice does not lead to simple and interpretable coefficients in percent, particularly when dealing with count data containing many zeros. The authors show that log-like transformations of the form  $\log(1 + Y)$  or  $\operatorname{arcsinh}(Y)$  introduce arbitrary scale dependence in the presence of

---

<sup>11</sup>The PTA is verified in the following part showing the event-study.

extensive margin effects. Specifically, the magnitude of treatment effects estimated using such transformations can be made arbitrarily large or small simply by rescaling the outcome variable - a property fundamentally at odds with the percentage interpretation often ascribed to these estimates. However, in our context, we need to acknowledge that a variation from zero to a certain value, or the opposite, can be viewed as a reflect of the extensive margin (a firm starts or stops producing patent), but also the intensive margin as a firm not producing patent a given year doesn't necessarily mean that they stopped innovating overall.

To address these methodological concerns while properly accounting for the count nature of patent data, we implement a multi-stage estimation strategy following the recommendations of the authors.

First, our main parameter of interest is  $\theta_{ATT\%}$ , which represents the proportional change in the average outcome for acquired firms:

$$\theta_{ATT\%} = \frac{E[Y_{it}(1)|D_i = 1, Post_t = 1] - E[Y_{it}(0)|D_i = 1, Post_t = 1]}{E[Y_{it}(0)|D_i = 1, Post_t = 1]}$$

This parameter is particularly appealing because it provides a scale-invariant measure of the treatment effect and has a natural percentage interpretation (in terms of the control group post period mean).  $\theta_{ATT\%}$  combines both intensive margin effects (changes in patent production for firms that continue to innovate) and extensive margin effects (firms that start or stop innovating altogether). While this does not distinguish between the two margins, we make the assumption for this estimation that zeros can be viewed as part of the intensive margin.

Our primary specification uses a Poisson QMLE event-study ([Wooldridge, 2023](#)):

$$Y_{it} = \exp \left( \lambda_t + D_i \beta_2 + \sum_{r \neq -1} D_i \times [RelativeTime_t = r] \beta_r^{ES} \right) \epsilon_{it} \quad (1)$$

where  $Y_{it}$  represents our patent-based measures for firm  $i$  in year  $t$ ,  $\lambda_t$  captures the year fixed effect,  $D_i$  is the treatment indicator, and  $\beta_r^{ES}$  captures the event study coefficients for each relative time period  $r$ . The Poisson specification naturally accommodates the count nature of the data while providing coefficients that have a clear interpretation as proportional effects. Thanks to the panel nature of our data, the pre-acquisition coefficients, if not significant allows to verify the parallel trend assumption, on which relies the estimation of the parameter.

Then, we examine how sensitive our results are to the presence of zeros, and thus the extensive margin, by implementing a calibrated log transformation approach. Specifically, we estimate effects using a transformation of the form:

$$m(y) = \begin{cases} \log(y) & \text{if } y > 0 \\ -x & \text{if } y = 0 \end{cases}$$

where  $x$  represents the value assigned to the extensive margin change. This allows us to explicitly specify how much weight we place on changes from zero to positive values relative to proportional changes among positive values. By varying  $x$ , we can assess how different valuations of the extensive margin affect our conclusions. Hence, our second specification uses a similar event-study approach :

$$m(Y_{it}) = \lambda_t + D_i \mathbf{B}_2 + \sum_{r \neq -1} D_i \times [\text{RelativeTime}_t = r] \mathbf{B}_r^{\text{ES}} + \epsilon_{it} \quad (2)$$

We complete this analysis by directly examining the impact of acquisitions on different margins of patent production. Specifically, we estimate a series of linear probability models for the likelihood of being closed in the post-acquisition period ( $P(\text{Closed}|D = 1)$ ) or the likelihood of stop producing any innovation ( $P(Y = 0|D = 1)$ ). We also do the estimates for the likelihood of producing different quantities of patents (1, 2, 3-5, 6-10, and more than 10 patents) in the post-acquisition period. It provides a simple interpretation as it allows to see the difference in percentages points between acquired and non acquired firms, and decompose our initial results.

# 5 Results

This section presents our empirical results, structured to provide evidence of a causal effect of acquisition on innovation. We first establish our baseline results using Poisson regression. We then delve deeper into the mechanisms of this effect by analyzing the extensive margin and heterogeneity using a calibrated log-transformation approach and linear probability models.

## 5.1 Main Results: Impact of Acquisitions on Innovation Output

Figure 7, which provides a visual check for the parallel trends assumption and the pattern of treatment effects over time <sup>12</sup>:




Figure 7: This figure plots event study coefficients from our baseline Poisson regression framework, showing the evolution of innovation output around the acquisition event. The y-axis represents the proportional change in innovation output relative to the pre-acquisition period (t-1). Vertical bars indicate 95% confidence intervals. The reference period is t-1, the year immediately preceding acquisition.

Our matching allowed to remove conveniently pre-trend differences between our treated and control group differences. It confirms that we’ve obtained a panel of firms that operates in the same sector, at the same sector and had prior to the event date the same innovative behavior. It justifies further the use of a framework similar to the differences-in-differences design. There is a significant albeit small drop at the time of the acquisition. It can be explained due to the nature of our data since patents production is measured annually, and not precisely at the daily level or monthly level. Hence, it takes a whole year before the

<sup>12</sup>The corresponding table is provided in the appendix.

effect can be visible in terms of magnitude. The years 1 to 3 in the post period indicates a stabilization of the reduction over time.

Table 7 quantifies these results by average the effect in the post-acquisition period<sup>13</sup>. It shows a significant decrease in patent output. Column (1) shows the estimates for equation (1). There is a decrease of 27.2% ( $\exp(-0.318) - 1 = -0.272$ ), on average, in innovation activity compared to the control group in the post period.

**The role of closure** A potential concern with this estimate is that it might reflect firm closures more than an actual reduction in innovative behavior. Even if that was the case, it would still be an interesting fact in itself as it would explained that the channel by which innovation is reduced post-acquisition is through closure. However, in Column (2), we control for firm closures by simply adding a dummy variable that equals 1 after the moment of a firm is closed.<sup>14</sup> We obtain a slightly smaller but still substantial, 23.4% decrease in patent output. The fact that this estimate remains significant implies that the decrease in innovation following acquisitions isn't solely driven by firms being absorbed by their acquirer, but also by surviving firms being less innovative.

	(1)	(2)
Post × Treatment	-0.318*** (0.029)	-0.267*** (0.028)
Proportional Effect	-0.272	-0.234
Controls	No	Yes
Number of Firms (Treated, Control)	4244	4244
Treated Group Mean (Pre, Post)	1.5	1.1
Control Group Mean (Pre, Post)	1.4	1.9

Table 7: Poisson Regression Estimates of Acquisition Impact

*Note:* This table presents Poisson regression estimates for the impact of acquisitions on innovation output. Column (1) shows the baseline regression without controls, while Column (2) adds firm-level controls. Standard errors (in parentheses) are clustered at the firm level. The mean innovation outputs by treatment group and time period are displayed in the bottom panel. Pre-period refers to the three years before acquisition while post-period covers the four years after, including the acquisition year. The proportional effect row shows the implied treatment effect calculated as  $\exp(\hat{\beta}) - 1$ . Statistical significance: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

<sup>13</sup>The average difference in the post-acquisition is the most economically significant figure as we wouldn't distinguish the effect between periods, but have a simple and clear answer as to whether or not being acquired is associated with less innovation.

<sup>14</sup>Naturally, a firm closure is highly correlated with the acquisition in itself. However, for the sake of our analysis, we hypothesize it as a random event.



## 5.2 Disentangling Extensive and Intensive Margin Effects

While our Poisson regression offers estimates of the average treatment effect on treated on innovation, it does not isolate how acquisitions affect the likelihood of a firm continuing to innovate (intensive margin) or the probability of firms stopping to innovate entirely (extensive margin). More specifically, the coefficient estimated gives us the average difference between the treated and control group in the post period, assuming there is no distinction between the intensive and the extensive margin. Even though zeros in the context of patents can be viewed as idiosyncratic (due to the random nature of the innovation process or the delays in the publication of patents), it could also reflect a stop in the innovation process. The specificity of our data makes it hard to distinguish, therefore we need to estimate the weight of the extensive margin in our initial results, or what could be seen as one. We first normalize the outcome so that 1 corresponds to the pre-treatment mean value of patent production, in a similar fashion as [Chen and Roth \(2024\)](#). Using equation (2), Figure 9 illustrates the variation in the estimated treatment effect on treated as we modify the weight placed on the extensive margin, controlled by parameter  $x$ :




Figure 8: This figure shows how the estimated treatment effect varies as we place different weights on extensive margin changes. The x-axis shows different values of the extensive margin parameter  $x$ , while the y-axis shows the estimated treatment effect under each calibration.

Figure 9 reports the coefficient for the event-study, and Table 8 the average coefficient for the post acquisition period.

When  $x$  is set to 0, effectively treating zero patent production as equivalent to the pre-treatment mean, we observe a treatment effect of -8.7 log points. This specification essentially

”shuts off” the extensive margin change between 0 and positive patent production, focusing solely on intensive margin changes. As we increase  $x$  to 1.0, valuing the extensive margin change at 100 log points, the magnitude increases substantially to -33.1 log points, suggesting that transitions to zero patent production play a significant role in the overall effect. Further increasing  $x$  to 3.0 and 10.0 leads to even larger estimated effects of -82.0 and -252.9 log points respectively<sup>15</sup>.

We can deduce from this analysis that the extensive margin (or at least the presence of zeros in the dataset) is what drives the difference between our treated and control group on average.

	(1)	(2)	(3)	(4)
Post $\times$ Treatment	-0.087*** (0.006)	-0.331*** (0.009)	-0.820*** (0.017)	-2.529*** (0.049)
Extensive Margin Value ( $x$ )	0.0	1.0	3.0	10.0

Table 8: Effect of Acquisition on Patent Production with Different Extensive Margin Values

*Note:* This table shows estimates of the treatment effect on the treated using  $m(Y)$  as the outcome, where  $m(Y)$  is defined to equal  $\log(Y)$  for  $Y > 0$  and  $-x$  for  $Y = 0$ . The outcome is normalized so that  $Y = 1$  corresponds to the pre-treatment mean value of the outcome. Thus, the treatment effect assigns a value of  $100x$  log points to an extensive margin change between 0 and the pre-treatment mean value of  $Y$ . The treatment effects are estimated using a difference-in-differences regression, using our transformed variable  $m(Y_{it})$ . Standard errors are clustered at the firm level. Statistical significance: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

**Understanding the effect on the extensive margin** Since the negative effect of acquisition is mostly due to the extensive margin, we try to provide a partial decomposition of our results through linear probability models. In our context, the extensive margin can be defined at two levels : a firm generally stop to innovate in the post acquisition period, or a firm is closed. Both could be used as an extensive definition in this context. Indeed, a firm can be acquired and forced to stop innovation in the post period, while not being closed, while a firm could be closed and therefore not produce innovation. Therefore, we estimate the effect on those two definitions of the extensive margin.

Table 9 summarizes the estimated effects:

<sup>15</sup>We don’t interpret log points as percentages following [Chen and Roth \(2024\)](#).

	(1)	(2)	(3)	(4)
	Closure		Stop Innovation	
	Post $\geq$ 0	Post $>$ 0	Post $\geq$ 0	Post $>$ 0
Post $\times$ Treatment	8.591***	8.643***	5.063***	10.358***
	(0.379)	(0.389)	(1.066)	(1.100)

Table 9: Linear Probability Model Estimates of Firm Closure and Stopping Innovation

*Note:* This table shows difference-in-differences estimates from linear probability models. All effects are measured in percentage points. Columns (1) and (2) estimate the effect on firm closure probability, while columns (3) and (4) estimate the effect on the probability of permanently stopping innovation in the post acquisition. Post  $\geq$  0 includes the acquisition year, while Post  $>$  0 excludes it. All effects are measured in percentage points. Standard errors (in parentheses) are clustered at the firm level. Statistical significance: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

These results show the prominent role of extensive margins in driving the decline in innovation following acquisitions. We define the post period in two ways. Indeed, due to timing, the year of acquisition may not necessarily capture the effect of the acquisition as patent produced for this specific year, but before the acquisition happened. Acquisitions increase the probability of firm closure by approximately 8.6 percentage points, regardless of the definition of the post period. However, the effect of acquisitions on innovation cessation shows a stronger temporal pattern. Acquisitions increase the probability of firms stopping innovation in the post-acquisition period by 5.1 percentage points when including the acquisition year and by 10.3 percentage points when excluding the acquisition year. It implies that acquisitions are more likely to stop innovation once they are acquired. Interestingly, even among firms that continue to exist post-acquisition, a notable proportion ceases to innovate.

**The distributional effect of being acquired** While the stopping effect on innovation can be associated to the extensive margin, we bypass the difficulty to handle zeros in our dataset by understanding the effect on being acquired on the production of a certain number of patents given the raw distribution. We define 5 categories of patent production in the totality of the post period (1, 2, 3 to 5, 6 to 10, and strictly more than 10). Table 10 provides estimates for the probabilities of producing patents in each category, conditional on being acquired:

	Post-Acquisition Patent Categories				
	1	2	3-5	6-10	>10
Treatment Effect	1.21*	-3.00***	-3.17***	-1.58**	-3.15***
	(0.69)	(0.66)	(0.74)	(0.62)	(0.59)

Table 10: Linear Probability Model Estimates of Patent Production Categories

*Note:* This table presents the differences-in-differences estimates for the linear probability model estimates of the effect of acquisition on the probability of producing patents in different categories in the post-acquisition period. All effects are measured in percentage points. Standard errors (in parentheses) are clustered at the firm level. Statistical significance: \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

The results show a small but statistically significant increase in the likelihood of producing exactly one patent after the acquisition. However, this small positive impact is outweighed by statistically significant decreases in the probability of producing patents in all other categories. These decreases range from 1.5 percentage points for 6-10 patents to over 3 percentage points for the higher categories. These findings indicate that the extensive margins drive main effects, but also that acquisitions cause a general shift toward low-intensity innovation by preventing higher levels of patent production, indicating a broader reduction in inventive efforts among acquired firms, and thus confirming our initial descriptive evidence.

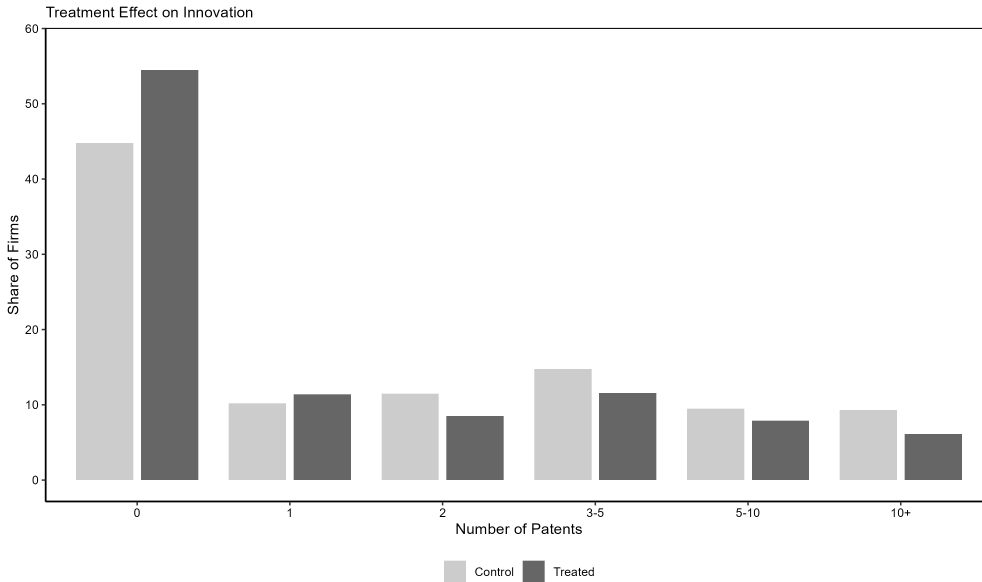


Figure 9: This figure shows the percentages of firms that produced a certain quantity of patents for both treated and control group.

## 6 Discussion

Recent theoretical literature suggests that acquisitions can be considered anti-competitive when the acquired firm had the potential to pursue innovation independently (Affeldt and Kesler, 2021b; Brutti and Rojas, 2022). The positive innovation trajectory observed in our control group indicates that, absent acquisition, most target firms would have likely maintained their innovative activities.

The interpretation of these results requires careful consideration of the *ex ante* innovation incentives faced by potential acquisition targets. A growing body of literature suggests that the prospect of acquisition itself may be a key driver of innovation. Einfeld (2024) estimates that prohibiting acquisitions could reduce start-up market entry by 8% to 20% across different markets, potentially diminishing both competition and innovation. Furthermore, incumbents' acquisition strategies can significantly influence the innovation decisions of new entrants, particularly in ways that enhance overall innovation as start-ups align their R&D efforts with incumbents' strategic needs (Warg, 2021; Wang, 2018). In practice, this suggests that even if acquired firms have, on average, contributed more to innovation (either independently or through their contributions to their acquirers) than non-acquired firms for a long period before their acquisitions, the observed overall decrease post-acquisition does not necessarily imply reduced total innovation contribution. Moreover, if part of the motivation for acquired firms to innovate was the prospect of being acquired, then some incremental innovations may have not existed, potentially leading to less overall innovation.

## 7 Conclusion

This paper provides evidence supporting the notion that acquisitions lead to a decrease in innovation for acquired firms, either through closure or reduced innovative output. Our results, given our counterfactual, suggest that the level of innovative output may have remained stable or even increased had the acquisition not occurred. This finding may imply that acquisitions can have a negative but moderate effect on innovation, except for cases where there is technological overlap. However, we believe this interpretation to be limited, as it does not fully account for the evolution of exit strategies and the relationship between acquisition and innovation for incumbents. If a firm innovates with the aim of being acquired, it implies that the higher level of innovation observed *ex ante* occurred due to the prospect of acquisition. In this context, innovation serves as a distinguishing factor and potentially as a strategy to cater to potential acquirers.

These findings have important implications for competition policy and merger control. For instance, the European Commission's (EC) evaluation process for mergers and acquisi-

tions (M&As) could benefit from considering some aspects linked to this paper, and all competition regulators. The EC should assess whether an acquisition might negatively affect the overall innovation dynamics of the market, including the potential for “killer acquisitions” that could hamper innovation from new entrants, as well as market contestability and the possibility of having “innovation spaces”<sup>16</sup>. Our research suggests that regulators could integrate into its reasoning the idea that while innovation may decrease post-acquisition, this decrease should be contextualized. The regulator should consider whether the acquired firm’s initial higher level of innovation was linked to the prospect of being acquired, or if other factors were at play. When examining a M&A, the regulator should take into account how the acquired firm was incentivized initially to innovate and assess whether the potential decrease in innovation *ex post* is negative or not at the sectoral level<sup>17</sup>.

To provide a more comprehensive answer to our initial question, further research could focus on how acquisitions may or may not be beneficial to innovation overall, with particular emphasis on the impact on the acquirer. Future work may also want to inquire about how the prospect of being acquired specifically incentivizes innovation for new entrants in a given market.

---

<sup>16</sup>The concept of “innovation spaces” was introduced by the European Commission in the context of the Dow/DuPont merger case. It refers to the areas of research and development where firms compete to innovate. They measure it using patent portfolio and patent overlap.

<sup>17</sup>We do not delve into this topic in the present article, as we believe the methodology required to make such an evaluation is beyond the scope of the current study.

## References

- Pauline Affeldt and Reinhold Kesler. Big Tech Acquisitions — Towards Empirical Evidence. *Journal of European Competition Law & Practice*, 12(6):471–478, June 2021a. ISSN 2041-7764, 2041-7772. doi: 10.1093/jeclap/lpab025. URL <https://academic.oup.com/jeclap/article/12/6/471/6232342>.
- Pauline Affeldt and Reinhold Kesler. Competitors’ Reactions to Big Tech Acquisitions: Evidence from Mobile Apps, December 2021b. URL <https://papers.ssrn.com/abstract=3998875>.
- Kenneth Arrow. Economic Welfare and the Allocation of Resources for Invention. *NBER Chapters*, pages 609–626, 1962. URL <https://ideas.repec.org/h/nbr/nberch/2144.html>. Publisher: National Bureau of Economic Research, Inc.
- Michele Bisceglia, Jorge Padilla, Joe Perkins, and Salvatore Piccolo. Optimal Exit Policy with Uncertain Demand, July 2023. URL <https://papers.ssrn.com/abstract=3889226>.
- Bruce A. Blonigen, Lionel Fontagné, Nicholas Sly, and Farid Toubal. Cherries for sale: The incidence and timing of cross-border M&A. Université Paris1 Panthéon-Sorbonne (Post-Print and Working Papers), December 2014. URL <https://ideas.repec.org/p/hal/cesptp/hal-01299761.html>.
- Zelda Brutti and Luis E. Rojas. M&A and early investment decisions by digital platforms. *Journal of Industrial and Business Economics*, 49(3):509–543, September 2022. ISSN 1972-4977. doi: 10.1007/s40812-022-00223-3. URL <https://doi.org/10.1007/s40812-022-00223-3>.
- Luís Cabral. Merger policy in digital industries. *Information Economics and Policy*, 54: 100866, March 2021. ISSN 01676245. doi: 10.1016/j.infoecopol.2020.100866. URL <https://linkinghub.elsevier.com/retrieve/pii/S0167624519302537>.
- Tyrone M. Carlin, Nigel Finch, and Guy Ford. A Deal Too Far: The Case of the Killer Acquisition. In Greg N. Gregoriou and Karyn L. Neuhauser, editors, *Mergers and Acquisitions: Current Issues*, pages 234–248. Palgrave Macmillan UK, London, 2007. ISBN 978-0-230-58968-1. doi: 10.1057/9780230589681\_15. URL [https://doi.org/10.1057/9780230589681\\_15](https://doi.org/10.1057/9780230589681_15).
- Jiafeng Chen and Jonathan Roth. Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics*, 139(2):891–936, 2024.

- Clayton M. Christensen. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Management of Innovation and Change. Harvard Business School Press, Boston, MA, 1997. ISBN 978-0-87584-585-2.
- Colleen Cunningham, Florian Ederer, and Song Ma. Killer Acquisitions. *Journal of Political Economy*, 129(3):649–702, March 2021. ISSN 0022-3808. doi: 10.1086/712506. URL <https://www.journals.uchicago.edu/doi/10.1086/712506>. Publisher: The University of Chicago Press.
- L. de Barys and A. Gautier. Big tech acquisitions and innovation: An empirical assessment. Working paper, CESIFO, 2024. Working Paper.
- Matthew Denes, Sabrina T. Howell, Filippo Mezzanotti, Xinxin Wang, and Ting Xu. Investor Tax Credits and Entrepreneurship: Evidence from U.S. States, August 2022. URL <https://papers.ssrn.com/abstract=3596342>.
- Luise Eisfeld. Entry and Acquisitions in Software Markets. Working Paper, 2024.
- Giulio Federico, Gregor Langus, and Tommaso M. Valletti. Horizontal mergers and product innovation. *ERN: Other Microeconomics: Production*, 2018. URL <https://api.semanticscholar.org/CorpusID:96450969>.
- Christian Fons-Rosen, Pau Roldan-Blanco, and Tom Schmitz. The Effects of Startup Acquisitions on Innovation and Economic Growth, February 2021. URL <https://papers.ssrn.com/abstract=3785485>.
- FTC. Non-HSR Reported Acquisitions by Select Technology Platforms, 2010-2019: An FTC Study. Technical report, Federal Trade Commission, 2021.
- Axel Gautier and Joe Lamesch. Mergers in the digital economy. *Information Economics and Policy*, 54:100890, 2021. ISSN 0167-6245. doi: <https://doi.org/10.1016/j.infoecopol.2020.100890>. URL <https://www.sciencedirect.com/science/article/pii/S0167624520301347>. Antitrust in the Digital Economy.
- Klaus Gugler, Florian Szücs, and Ulrich Wohak. Start-up Acquisitions, Venture Capital and Innovation: A Comparative Study of Google, Apple, Facebook, Amazon and Microsoft. Working Paper, 2023.
- Justus Haucap, Alexander Rasch, and Joel Stiebale. How mergers affect innovation: Theory and evidence. *International Journal of Industrial Organization*, 63:283–325, March 2019. ISSN 0167-7187. doi: 10.1016/j.ijindorg.2018.10.003. URL <https://www.sciencedirect.com/science/article/pii/S0167718717303685>.



- Matthew J. Higgins and Daniel Rodriguez. The outsourcing of R&D through acquisitions in the pharmaceutical industry. *Journal of Financial Economics*, 80(2):351–383, May 2006. ISSN 0304-405X. doi: 10.1016/j.jfineco.2005.04.004. URL <https://www.sciencedirect.com/science/article/pii/S0304405X05001807>.
- Brett Hollenbeck. Horizontal mergers and innovation in concentrated industries. *Quantitative Marketing and Economics*, 18(1):1–37, March 2020. ISSN 1573-711X. doi: 10.1007/s11129-019-09218-2. URL <https://doi.org/10.1007/s11129-019-09218-2>.
- Marc Ivaldi, Nicolas Petit, and Selcukhan Unekbas. Killer Acquisitions: Evidence from EC Merger Cases in Digital Industries. *SSRN Electronic Journal*, 2023. ISSN 1556-5068. doi: 10.2139/ssrn.4407333. URL <https://www.ssrn.com/abstract=4407333>.
- Bruno Jullien and Yassine Lefouili. Horizontal Mergers and Innovation, August 2018. URL <https://papers.ssrn.com/abstract=3135177>.
- Sai Krishna Kamepalli, Raghuram G. Rajan, and Luigi Zingales. Kill Zone, February 2021. URL <https://papers.ssrn.com/abstract=3555915>.
- Mark A. Lemley and Andrew McCreary. Exit Strategy, December 2019. URL <https://papers.ssrn.com/abstract=3506919>.
- Igor Letina, Armin Schmutzler, and Regina Seibel. Killer Acquisitions and Beyond: Policy Effects on Innovation Strategies, October 2021. URL <https://papers.ssrn.com/abstract=3673150>.
- Ben Mermelstein, Volker Nocke, Mark A. Satterthwaite, and Michael D. Whinston. Internal versus External Growth in Industries with Scale Economies: A Computational Model of Optimal Merger Policy, April 2014. URL <https://www.nber.org/papers/w20051>.
- Gordon M. Phillips and Alexei Zhdanov. R and D and the Incentives from Merger and Acquisition Activity. *The Review of Financial Studies*, 26(1):34–78, 01 2013. ISSN 0893-9454. doi: 10.1093/rfs/hhs109. URL <https://doi.org/10.1093/rfs/hhs109>.
- Felix Poege. Competition and Innovation: The Breakup of IG Farben, August 2022. URL <https://papers.ssrn.com/abstract=4199672>.
- Tiago S. Prado and Johannes M. Bauer. Big Tech Platform Acquisitions of Start-ups and Venture Capital Funding for Innovation, March 2022. URL <https://papers.ssrn.com/abstract=3787127>.

- Eric Rasmusen. Entry for Buyout. *Journal of Industrial Economics*, 36(3):281–99, 1988. URL [https://econpapers.repec.org/article/blajindec/v\\_3a36\\_3ay\\_3a1988\\_3ai\\_3a3\\_3ap\\_3a281-99.htm](https://econpapers.repec.org/article/blajindec/v_3a36_3ay_3a1988_3ai_3a3_3ap_3a281-99.htm). Publisher: Wiley Blackwell.
- Joseph A. Schumpeter. *Capitalism, Socialism, and Democracy*. Harper & Brothers, New York, 1 edition, 1942.
- Amit Seru. Firm boundaries matter: Evidence from conglomerates and rd activity. *Journal of Financial Economics*, 111(2):381–405, 2014. ISSN 0304-405X. doi: <https://doi.org/10.1016/j.jfineco.2013.11.001>. URL <https://www.sciencedirect.com/science/article/pii/S0304405X1300278X>.
- D. Daniel Sokol. Merger Law for Biotech and Killer Acquisitions, September 2020. URL <https://papers.ssrn.com/abstract=3658337>.
- Mariagrazia Squicciarini, H el ene Dernis, and Chiara Criscuolo. Measuring patent quality: Indicators of technological and economic value. *OECD Science, Technology and Industry Working Papers*, 2013.
- Gianluca Tarasconi and Carlo Menon. Matching crunchbase with patent data, 2017. URL <https://www.oecd-ilibrary.org/content/paper/15f967fa-en>.
- Xinxin Wang. Catering Innovation: Entrepreneurship and the Acquisition Market, July 2018. URL <https://papers.ssrn.com/abstract=3247274>.
- Katarina Warg. The Acquisition Option and Startup Innovations., 2021.
- Martin Watzinger, Thomas A. Fackler, Markus Nagler, and Monika Schnitzer. How Antitrust Enforcement Can Spur Innovation: Bell Labs and the 1956 Consent Decree. *American Economic Journal: Economic Policy*, 12(4):328–359, November 2020. ISSN 1945-7731. doi: 10.1257/pol.20190086. URL <https://www.aeaweb.org/articles?id=10.1257/pol.20190086>.
- Jeffrey M Wooldridge. Simple approaches to nonlinear difference-in-differences with panel data. *The Econometrics Journal*, 26(3):C31–C66, 08 2023. ISSN 1368-4221. doi: 10.1093/ectj/utad016. URL <https://doi.org/10.1093/ectj/utad016>.
- John M. Yun. Discriminatory Antitrust in the Realm of Potential and Nascent Competition. George Mason Law & Economics Research Paper No. 22-04., 2022.

## A.1 Appendix

### A.1.1 Methodologies for Data and Variables

#### A.1.1.1 Embeddings, and Similarity for Industry Retrieval

**Using embeddings from Mistral AI** To define relevant markets for each firm, we leverage the embedding generation capability of the Mistral AI Large Language Model (LLM) embedding generator. We choose this method and this particular model for two specific reasons. First, traditional methodologies for computing similarity between two texts (for example, product descriptions or patent descriptions) often relied on Latent Semantic Analysis (LSA). However, LSA is based on the statistical co-occurrence of words and does not necessarily take into account the context of the sentence and its semantic meaning.

For instance, the two sentences “Our company develops software solutions for project management” and “Our firm designs planning tools for the construction industry” may be deemed quite similar due to the co-occurrence of some words (solutions, tools). The first sentence, however, clearly positions the company in the IT sector, while the other places it in the construction industry. This limitation is overcome by the generation of embeddings. These embeddings, denoted as  $e_i = f_{\text{Mistral}}(d_i)$ , are created from a concatenation of the company name, Crunchbase industries, and full description. When the full description isn’t available, we use the short description.

**NAICS Code Assignment and Sectoral similarity** For NAICS code assignment, we first compute reference embeddings for each 6-digit NAICS category using the same Mistral AI model:  $e_{\text{NAICS}_j} = f_{\text{Mistral}}(\text{description}_{\text{NAICS}_j})$ . We then calculate the cosine similarity between each company’s embedding and these NAICS embeddings, assigning the two codes with the highest similarity scores. It is often the case that a firm may be classified into several NAICS codes. We believe that using two NAICS codes to classify a firm is sufficient to characterize it with enough precision at the sectoral level. Unfortunately, we did not have access to a complete database containing a list of firms and their actual NAICS codes (as registered administratively) to test the quality of our predictions. However, manual testing on a small but relevant sample has led us to believe that we could match it with sufficient precision. Further research and potentially open-source databases on classifying firms may be very helpful to practitioners.

We employ a similar approach to estimate the similarity between two companies, typically a targeted firm and its acquirer. The cosine similarity of their embeddings is calculated as:  $\text{sim}(e_i, e_j) = \frac{e_i \cdot e_j}{\|e_i\| \|e_j\|}$ . Companies are considered to be in the same sector if their similarity exceeds a threshold of 0.7. While this threshold is somewhat arbitrary, as there is

no universally accepted value, it was chosen based on empirical observations. This value effectively captured the similarity between firms that we identified as similar through human understanding of their descriptions. It represents the lowest threshold at which we could consistently detect this similarity. Although we cannot claim that this value is optimal, it provides a reliable indication of sectoral similarity, especially when combined with other complementary measures to create a comprehensive similarity indicator.

#### **A.1.1.2 Building OpenPat**

For the purposes of this article and future research, we required a comprehensive dataset on patents. While the United States Patent Office (USPTO) provides valuable data, it has several limitations: it is restricted to one jurisdiction, and some American startups, like Basepaws, file patents in other jurisdictions rather than in the US. Moreover, existing comprehensive patent databases are either prohibitively expensive (e.g., PatStat) or lack user-friendly interfaces for data extraction (Google Patents or Lens.org only provide search interfaces without structured datasets of all patents).

**Data Collection** We utilized Lens.org, which allows downloading structured datasets containing unique IDs, jurisdictions, inventor names, applicant entities, owners, citation counts, number of patents cited, simple family weight, titles, and abstracts. However, downloads are limited to 50,000 patents per search (based on specific time periods and category filters). This necessitated 6,150 manual downloads<sup>18</sup>.

**Data Processing** After gathering the data, we created a variable called “Final Selection” to identify the correct entity and serve as our reference name. We removed inventor names from the applicant variable by “subtracting” the inventors from the applicants, focusing solely on legal entities.

To choose the correct names for the Final Selection, we followed these rules:

1. If there are no owners, the Final Selection is the Applicant.
2. We prefer owners over applicants by default.
3. For cases with one applicant and multiple owners (due to acquisitions), we choose the owner most similar to the applicant name based on string similarity, as we are interested in the firm that published the patent.

---

<sup>18</sup>This process is time-consuming as the data is sent via email, not always successfully due to possible interface limitations, and may take several minutes or more to receive.

4. We separated our Final Selection to associate all potential companies in our data with the correct owners.

**Standardization and Matching** To standardize and accurately match the Final Selection with Crunchbase data, we utilized the IPQwery Crunchbase Extension, which identifies firms owning patents in a dataset similar to PatStat. This provided crucial indicators for likely matches in our dataset. The Crunchbase dataset contains approximately 90,000 firms with patents.

We performed the following steps:

1. Cleaned and harmonized names by removing stop words and legal terms (e.g., LLC).
2. Conducted exact matching to link as many firms and Final Selections as possible, especially those with unique names.
3. Implemented fuzzy matching, incorporating variables such as firm size, founding date, and IPO status from Crunchbase, and patent counts from our OpenPat database.
4. Used a threshold of 0.8 for fuzzy matching, balancing accuracy and match rate.
5. Employed Mistral AI embeddings to facilitate name matching.

**Results** We successfully matched around 161,000 Final Selections in our OpenPat data, which is comparable to similar studies ([Tarasconi and Menon, 2017](#)), taking into account the fact that there is a 7 year gap between ours and the OECD one. In future research, we aim to create a more comprehensive OpenPat Database with improved harmonization of the Final Selection, making it more compatible with Crunchbase data.

## A.1.2 Additional Tables

### A.1.2.1 BATX, GAFAM, NATU

Table 11

---

Characteristic	BATX	GAFAM	NATU
<i>Acquired firms characteristics</i>			
Size (median)	11-50	11-50	11-50
Relative age (median)	7	5	4
Patent owners (%)	2.9	6	6
Patent applicants (%)	4.8	18.4	12
VC-backed (%)	83.8	94.9	90.3
Number of rounds (median)	2	2	2
Created after 2000 (%)	87.1	76.7	78.7
<i>Compared to acquirer</i>			
Sectoral similarity (%)	53.8	47.8	64
Acquiree is younger (%)	86.1	94.3	80.9
Acquiree is smaller (%)	93.3	95.1	97.9
Patent overlap among owners (%)	50	62.6	16.8
Total observations	104	761	50

---

### A.1.3 Event-study Poisson

	(1)	(2)
$t - 3$	-0.023 (0.063)	-0.024 (0.063)
$t - 2$	-0.055 (0.039)	-0.056 (0.039)
$t + 0$	-0.047 (0.039)	-0.026 (0.040)
$t + 1$	-0.223*** (0.044)	-0.194*** (0.045)
$t + 2$	-0.270*** (0.058)	-0.229*** (0.060)
$t + 3$	-0.287*** (0.060)	-0.240*** (0.062)
Covariates	N	Y

Table 12: Event Study Estimates

Note: This table shows the event study coefficients for each period relative to the acquisition. The coefficients represent proportional treatment effects. Period  $t - 1$  is omitted as the reference period. Stars indicate significance levels (\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ ). Standard errors are clustered at the firm level.