HEC LIÈGE
UER ÉCONOMIE

Big Tech Acquisitions and Innovation Strategies

Dissertation submitted
in fulfilment of the requirements for the degree
of Ph.D. in Economics and Management

Laureen DE BARSY

Advisor

Axel GAUTIER	HEC Liège, Université de Liège

Jury

Patrice BOUGETTE	Université Côte d'Azur
Malka GUILLOT	HEC Liège, Université de Liège
Valerio STERZI	Bordeaux School of Economics
Barnabé WALHEER	HEC Liège, Université de Liège

2024

Acknowledgments

I am very grateful to the many people who made this dissertation possible. First, I would like to thank my supervisor, Axel Gautier, for his support and encouragement over the last three years. His trust and kindness meant a lot to me. I have learned a lot from him as a supervisor and co-author, and his genuine excitement for his work has been truly inspiring.

I also want to thank the members of my thesis committee. Thanks to Malka Guillot, on whom I felt I could always count. Her direct and constructive advice has been invaluable. Thanks to Valerio Sterzi, for his generosity in sharing knowledge about patents, and for providing me with a second academic home in Bordeaux. I am also very grateful to Patrice Bougette and Barnabé Walheer, for their precious feedback and encouragement.

I would also like to thank Maren Ulm, Sylvain Dejean and Joseph Tharakan for their insightful comments and suggestions on my thesis chapters. My work has also benefited from conversations with Luise Eisfeld, Antonio Estache, Michele Bisceglia, Giovanni Morzenti and Cristina Rujan, to whom I am very thankful. I also want to acknowledge the valuable support of Guadalupe Amésquita Palacios at every step of the doctoral program.

My colleagues, both past and present, have been instrumental in improving and completing this work. Thanks to all the teaching assistants, PhD students and researchers of the UER Economie, for the warm welcome in their team. The friendly atmosphere they created made a significant difference during this often solitary experience. Special thanks to Hugues, an amazing office mate, to Kelly and Xavier, two pillars of the scientific staff, to Laure, a great Python coach, and to the always cheerful Sousso. Thanks also to Franziska Valder and Enrico Camarda for their mentorship and friendship. I am also grateful to Andreas Bovin for helping me retrieve important data for my third chapter.

This PhD thesis likely would never have come to fruition without the support of my family and friends. Thanks to Simon for his love, for being such a good audience to my conference rehearsals, for his delicious meals, and for supporting me during challenging times. I am also fortunate to have received love and support from my close friends, who provided plenty of distractions and joy away from work. Special thanks to Federica and Laurens, for being a constant source of encouragement. Lastly, my deepest gratitude goes to my parents, for believing in me and giving me every opportunity to pursue what inspires me.

Laureen
Liège, September 19, 2024.

Contents

List of Figures

List of Tables

General Introduction: Regulatory Framework and Market Dynamics

If none of it's of interest to you, you'd be the first
—Bo Burnham, "Welcome to the Internet"

0.1 The Rise of Big Tech Platforms

Since its introduction to the public in the mid-1990s, the Internet has evolved into a global digital economy of unprecedented scale. The digital landscape has reshaped how businesses operate and interact with consumers, and the economy is today heavily dependent on digital and internet technologies. With more than 5 billion users, this new economy is characterized by a rapid pace of technological and conceptual innovation that makes it a highly competitive arena (Schimmer, Mueller-Stewens, and Sponland 2010).

Google (Alphabet), Apple, Facebook (Meta), Amazon and Microsoft (often grouped under the labels Big Tech, GAFAM or tech giants) have long been the most successful companies in this arena.[1] Their strategic actions have not only shaped the information

[1] Today, Big Tech platforms are ranked in the top 10 most visited multi-platform U.S. web properties (Statista 2024).

economy as we know it today but also serve as prime examples of how profitable market positions can be achieved online (Schimmer, Mueller-Stewens, and Sponland 2010). Globalisation and technological change (e.g. increased automation and digitisation) have created opportunities for these highly productive firms to grow at the expense of less efficient rivals (De Loecker, Eeckhout, and Unger 2020, Valletti and Zenger 2019), who may not have been able to exploit the potential of digital technologies (Criscuolo 2019). The slow diffusion of digital technologies and the costly, time-consuming investments required in complementary intangibles — such as data, proprietary software, and human and organisational capital — have become crucial competitive assets that reinforce market leaders' positions (Brynjolfsson, D. Rock, and Syverson 2019, Crouzet and Eberly 2019, Jovanovic and Rousseau 2005).

In the recent years, the persistence and strengthening of the tech giants' market positions have been raising concerns.[2] In the short run, the success of these firms has provided significant benefits to users. However, in the long run, the growing concentration of digital markets seems to reduce business dynamism (Criscuolo 2019, Valletti and Zenger 2019). In response, reforms to the legal framework regulating digital industries are being discussed. In September 2022, the European Parliament and the Council adopted a regulatory tool aiming to ensure a greater degree of competition in the European market for digital services: the "Digital Markets Act". This legal framework regulates the activities of large digital platforms through a series of obligations (e.g. interoperability) and prohibitions (e.g. self-preferencing).

Big Tech's strong market position also has implications for the impact of their acquisitions strategies: "Mergers are more prone to be problematic when the pricing power of merging firms is already large to begin with" (Valletti and Zenger 2019: 49). The high pace at which these incumbent platforms acquire smaller players of the sector is often perceived as a threat to healthy competition. As of today, Big Tech platforms have bought more than 800 companies (see Figure 1). Some of these acquisitions attract a lot of media interest because of the notoriety of the acquired company and the sums spent

[2]See for instance ACCC (2019), Crémer, Montjoye, and Schweitzer (2019), Furman et al. (2019), or Scott Morton et al. (2019).

by the buyer - e.g. the acquisition of WhatsApp by Facebook in 2014 for $22 billion, or LinkedIn by Microsoft in 2016 for $27 billion, although few of them have been reviewed by a competition authority[3]. Even when they are reviewed, competition authorities are often unable to demonstrate a likely consumer harm.[4]

Figure 1: Big Tech acquisitions over time

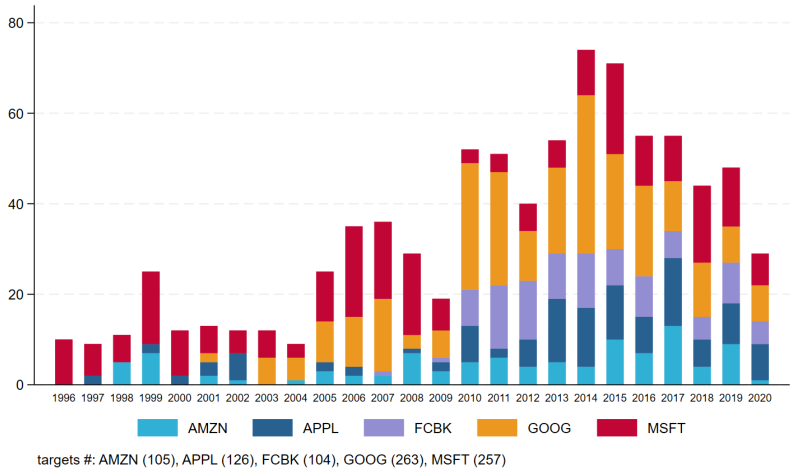

Source: Author's compilation based on Standard Poor's CapIQ (2022), Parker et al. (2021), Gautier and Lamesch (2021), and the USPTO Patent Assignment Dataset (2022).

Note: The decrease in the number of Big Tech acquisitions since 2015 should be weighed against the fact that acquisitions after that date are on average associated with bigger target firms (in terms of funding amounts and number of employees, Orbis Global data and author's computations).

[3] Around 97% of acquisitions in the technology sector have reportedly never been subject to scrutiny by a competition authority (Kwoka and Valletti 2021). Many digital mergers are not reviewed because they do not meet the legal threshold for intervention in terms of the individual turnover of the acquired start-up (See Article 1(2) of the European Merger Regulation, EC/139/2004.). In response, since March 2021, the European Commission allows Member States to refer to it the examination of transactions that do not meet the turnover thresholds if the turnover does not reflect the actual or future competitive potential of at least one of the merging parties (Guidance on the application of the referral mechanism set out in Article 22 of the Merger Regulation 2021/C 113/01, 2.2§19.).

[4] There are two recent exceptions from the UK Competition and Markets Authority (CMA): Microsoft's 2023 blocked deal to buy Activision, a leading video games publisher, and Giphy, one of the largest distributors of gifs on the Internet, that Facebook was reordered to sell in 2022.

0.2 Evaluating Big Tech Acquisitions

Competition authorities are in charge of controlling a market that is becoming more complex and opaque every day, and over which digital platforms have an advantage in terms of access to information thanks to the data they collect on their users (Parker, Petropoulos, and Van Alstyne 2021). This difficulty in accurately assessing the anticompetitive potential of Big Tech acquisitions can be partially explained by some specificities of digital markets: large returns to scale (i.e. the cost of producing a digital service is much less than proportional to the number of users served), network externalities (i.e. the larger the number of users of a digital service, the greater the value of that service), and the key role of data (Crémer, Montjoye, and Schweitzer 2019).

By giving incumbents a strong competitive advantage, these characteristics of digital markets foster their concentration. It is crucial that the methodology employed to review digital mergers accounts for these specificities, as failing to do so would lead to wrong decisions. For instance, after authorising the acquisition of Whatsapp by Facebook,[5] the Commission recognized that it had overlooked the key role played by network externalities. Since the Facebook/Whatsapp merger, users looking for a messaging service used by a large number of their contacts found themselves locked into the Facebook ecosystem. As a result, Facebook's user base continued to grow and became increasingly valuable to companies looking to target their advertising. Now essential, Facebook can afford to offer advertisers ever more expensive advertising tools without seeing their demand diminish.

This price effect is well understood by economic theory; in a static environment absent efficiencies and synergies, a horizontal merger relaxes a competitive pressure, which leads to higher prices, restricted output and a lower consumer surplus. But mergers also have an effect on innovation, and thus on future prices and products quality. And innovation is key in the digital world; the tech giants are spending billions in R&D and

[5]The European Commission authorised the Facebook/Whatsapp merger because it considered that there were sufficient alternative messaging services (See Commission decision of 3 October 2014 in Case M.7217 – Facebook/WhatsApp).

many of the firms they acquired are young and innovative startups that often develop new technologies.

A merger between two firms can have positive effects on innovation, and this could be used as an argument in the "balance of harm" approach of competition authorities. The EU and US reviewing agencies consider the potential innovative benefits of a merger in the context of "efficiencies" (Esteva Mosso 2018). For instance, in "TomTom/Tele Atlas" (2008), the European Commission recognised that the merger between a navigation systems provider and a digital maps developer would allow to deliver "better maps - faster". At the same time, a merger might also harm firms' incentives to innovate, as illustrated with the "Dow/DuPont" (2017) merger case in which the European Commission expressed concerns that the merger would have reduced innovation. In order to complete the merger, the parties agreed to divest assets in overlapping markets to preserve the industry's incentives to innovate. In some cases, mergers are even used to kill innovative products that threaten those of the incumbents, as documented by Cunningham, Ederer, and Ma (2021) for the pharmaceutical industry. In the digital industry, there is a fear that the acquisition of start-ups by a dominant platform results in the strengthening of its dominance, a reduction of effective competition, and a loss of innovation (Motta and Peitz 2021).

In practice, the standard approach to assessing the effects of a merger on innovation is to define a counterfactual, i.e. how innovation would evolve absent the merger. But there are many channels through which Big Tech acquisitions impact innovation, and the associated predictions are not necessarily aligned. In addition, innovation is uncertain by nature. Answering the *ex-ante* question of how innovation would develop if the incumbent platform was not allowed to take over a start-up is thus very challenging (De Coninck and Muellern 2023), making it difficult to predict the effect of an acquisition on competition: "Uncertainty about what products the incumbent and the nascent competitor will actually offer in the future has a further consequence - uncertainty about the degree to which those products will actually compete." (Hemphill and Wu 2020: 1887–88). With this thesis, I propose to take an *ex-post* approach instead, and to study the observed innovative activity around Big Tech acquisitions.

0.3 Contributions of this Thesis

Innovation is difficult to measure and cannot be quantified as easily as a product's price. Because my goal is to track innovation activities around Big Tech acquisitions, I make use of patent data that allows me to capture the transferable components of Big Tech-acquired technologies.

Unlike products, that often change names, patented technologies can be tracked as they move across firms. In the first chapter of my thesis, I use this feature of the patent data to assess the impact of Big Tech acquisitions on the development of acquired technologies. But patents expire and technologies evolve. In contrast, the talent of the people who generated the innovation is an asset that can produce a continuous flow of future innovation. Therefore, my second chapter explores another strategy Big Tech uses to access innovative technologies; acquiring the talent behind these technologies. Finally, next to examining how much acquired technologies are developed and how productive the acquired inventors are, it is important to consider the nature of the innovation itself. Some innovation just marginally improves an existing technology, while some other completely revolutionises the state of the art. In the third chapter, I propose shifting the focus to the disruptive nature of Big Tech-acquired technologies.

My results suggest that, when buying an innovative start-up, Big Tech mainly acquires a *potential*, a capability to innovate, rather than a fully developed product. The development of the (existing) acquired technologies is often just a temporary phase. Instead, the tech giants tend to leverage the acquired inventors' potential to drive further innovation. Finally, I show that, by strategically acquiring innovative start-ups, Big Tech can disrupt markets in which it holds a weaker position.

Chapter 1

Tracking Technology Developments after a Big Tech Acquisition[1]

Abstract

In the past 20 years, large digital platforms have made many acquisitions, mainly young and innovative startups. Few of them have been reviewed by competition authorities and little is known about the evolution of the startups after they have been acquired. This paper intends to fill this gap by looking at the development of the technologies owned by the acquired firms. We focus on technologies protected by patents and we investigate whether an acquisition by a tech giant contributes to their development. For this analysis, we use citations to acquired patents as a proxy for the development of the acquired technology, and we identify where these citations originated. We show that acquisition temporarily increases the development effort on the acquired technology. Interestingly, the decline following this initial boost starts earlier for citations originating from the acquirer than for citations by the rest of the industry. We also find that the slow down in the acquirer's development effort is stronger for acquired inventions in large patent portfolios.

[1]This chapter is co-authored with Axel Gautier.

1.1 Introduction

One of the most notable transformations of our economy over the last 30 years is its move towards digitisation. Google (Alphabet), Apple, Facebook (Meta), Amazon and Microsoft (which are often grouped under the labels GAFAM or Big Tech) supported that transformation by bringing more and more social and economic activities to the online world. From almost non-existent in the early 1990s, these companies are now among the most valuable companies worldwide.

Being the primary gateways through which people use the Internet places Big Tech in a position of dominance in digital markets. In order to maintain quality services at reasonable prices, regulators and competition authorities must ensure that other market players can still enter digital markets and compete with these dominant firms. Among the many challenges that the digital economy poses in that regard (e.g. strong network effects, multi-sidedness, data-driven economies of scope, etc.), the role of mergers and acquisitions (M&A) by Big Tech is increasingly considered[2], especially given the very high rate at which these platforms acquire start-ups. In an interview on CNBC[3], Tim Cook, Apple's CEO, illustrated that: *"We acquire everything that we need that can fit and has a strategic purpose to it. And so we acquire a company on average, every two to three weeks."*.

Next to the expected effect on prices coming from the loss of a competitive pressure, competition authorities also consider the likely impact of Big Tech start-ups acquisitions on innovation. The existing evidence shows that the startup's products are often no longer developed after acquisition. Gautier and Lamesch (2021) found that 60% of the products of firms acquired by the big techs are no longer supplied, maintained or upgraded after acquisition. Affeldt and Kesler (2021) focus on mergers involving 'apps' and they document that half of those apps were discontinued after an acquisition by a tech giant. Eisfeld (2022) studies startup acquisition in the software industry and finds that 57% of the acquired products have been discontinued under their original brand name

[2] See for instance Argentesi et al. 2021; Crémer, Montjoye, and Schweitzer 2019; "Stigler Report" 2019.
[3] Berkshire Hathaway's annual shareholder meeting, interview by Becky Quick on CNBC in 2019.

after acquisition.[4] However, product discontinuation does not mean that the acquired technology is no longer used, as it could continue to exist under a new brand name, be integrated in a new product or more generally in the acquirer's ecosystem. As a matter of fact, little is known about the development of technologies after acquisition. This paper intends to fill this gap.

After acquisition, the target becomes part of the tech giant. Engineers, research labs, projects and products are transferred to the acquirer and integrated in its ecosystem. To assess the impact of big tech acquisitions on innovation, and instead of tracking product-level development, this paper focuses on the projects' underlying technology, materialised by patents. By tracking patents as they move across firms, we are able to identify whether a technology continues to be developed after acquisition. More specifically, the patent system is such that, when some inventors build on an existing technology, they must cite the patents protecting that technology. This implies that the development of a technology is materialised by citations that are made to the patents protecting it. The number of citations made by the acquirer itself thus reflects the intentions of the acquirer towards this technology; a technology that it wants to develop will receive more citations than a technology that is destined to stagnate. We can therefore use the citations made by the acquirer as a proxy for its innovation effort to develop the acquired technology. A higher (lower) research effort translates into more (less) citations to the acquired patents. We thus intend to assess the impact of acquisition on the development of acquired technologies, as proxied by citations to the acquired patents.

For our analysis, we construct a sample of firms acquired by Big Tech since 1996 and we identified those that have filed some patents prior to their acquisition. Some acquired firms do not own patents either because they did not develop technologies or because they did not patent the technologies they developed. Not patenting an invention could be a strategic decision (e.g. firms that do not wish to disclose information could prefer secrecy over patenting, Arundel 2001), but it could also simply derive from a low probability of imitation, high costs of patenting (e.g. administrative costs and renewal

[4] Product discontinuation is particularly a concern when the target is small (Gautier and Lamesch 2021).

fees), the length of the grant procedure,[5] or from the conditions for patentability not being met (Belleflamme and Peitz 2015). In our sample, 29% of the acquired firms have a patented technology at the time of acquisition. This represents 76% of the 133 biggest acquired firms (i.e. with a total funding above $10 million).[6]

Next, we retrieve all the citations made by the acquirer to the acquired patents. We use the evolution of these citations as a proxy for the improvements by the acquirer to the acquired technology. By exploiting the time series nature of our data, we develop a methodology to identify the dynamic effect of acquisition on Big Tech citations to acquired patents. Life-cycle and business-cycle trends in the evolution of Big Tech citations are captured by controlling for the patent age and the date at which the citation was made. In a first model specification, the short-term impact of acquisition is identified from the sharp breaks in citations trajectories immediately following acquisition. Second, using propensity scores, we introduce a control group of non-acquired patents that are comparable to acquired patents. We then compare the remaining time trends in Big Tech citations to acquired patents with respect to comparable non-acquired patents using a difference-in-differences design.

In our analysis, we consider the number of citations to acquired patents during a period of 4 years around acquisition. Our empirical analysis shows that acquisition seems to first give a boost to the development of the acquired technology as citations increase directly after the merger. But, after 1.5 year, the developments made by Big Tech to acquired technologies start slowing down. We observe that citations by the acquirer follow an inverse U-shaped curve and this result is robust to many specifications that we tested. This suggests that the boost in the development of the technology by its acquirer fades away in the long run.

A possible explanation for the inverse U-shaped curve in citations after acquisition is that the acquired technologies are close to maturity and need few developing steps before

[5] US patents take approximately 32 months from their filing date to be granted (as computed based on the 'grant lag' from the OECD Patent Quality Indicators database, July 2021).

[6] Based on funding data retrieved from Crunchbase.

being commercialised. In such a case, we should observe a similar citation pattern in the rest of the industry. To test for this hypothesis, we look at the evolution of citations by the other firms in the industry, i.e. citations by the non-acquiring firms. Our analysis does not corroborate this technology maturity hypothesis as we observe that the rest of the industry keeps further developing these technologies up to 2.5 years after their acquisition. On this basis, we conclude that the improvement potential of the technology has not been exhausted after acquisition, so technology maturity alone is unlikely to explain Big Tech's declining interest for the development of acquired technologies.

Finally, we test whether our baseline results on the technology development by its acquirer vary across acquired patents portfolios of different sizes. We find that the boost in the acquirer's citations just after acquisition is stronger and more persistent for patents belonging to relatively small patent portfolios. For patents belonging to large portfolios, we do not observe a significant boost in citations just after acquisition and the slow down after 1.5 year is more pronounced. A potential explanation for this result is that, for small targets, the acquisition decision is largely driven by a specific technology, whereas large targets may own many patents that are of little interest to the acquirer.

Related literature

Our paper is related to the literature on mergers and innovation. This literature studies the impact of a merger on the innovation by the merging entity, the competitors and the acquired company.

The start-up innovative effort can first be impacted through the possibility of buyout. In case it does not manage to bring its project to the market, a start-up might want to secure the outside option of being acquired by a bigger firm. The prospect of acquisition can boost the innovative effort of the start-up because the acquisition rents increase the expected profit from innovation (Cabral 2021, Motta and Peitz 2021). However, while doing so, the start-up might strategically distort the direction of its innovation in order to maximise the acquisition rents (Dijk, José L Moraga-González, and Motchenkova 2021, Katz 2021).

Mergers might also impact innovation by the acquirer's competitors, actual or potential. When firms are competing in innovation, a merger has an impact on the innovation effort of the outsiders to the merger. Federico, Langus, and Valletti (2018) show that a merger reduces the innovation effort by the merged entity but increases the research effort of the competitors (i.e. research efforts are strategic substitutes). Innovation by actual competitors might be hindered when startups that could have enabled them to catch up technologically are bought by the leading platform (Bryan and Hovenkamp 2020a).

Empirically, the effect of digital M&As on innovation by competitors of the merging entity has been tackled in a recent study by Affeldt and Kesler (2021). These authors study Big Tech acquisitions in the Google Play Store. They find that, after the acquisition of an app by a tech giant, competing apps are less likely to be invented or updated and developers shift their innovation effort to non-competing apps. Koski, Kässi, and Braesemann (2023) and Eisfeld (2022) study the impact of mergers on potential competitors. Koski, Kässi, and Braesemann (2023) provide evidence that mergers decrease entry. Eisfeld (2022) has more nuanced results; she shows that a more stringent merger policy would reduce entry, as buyout is one of the main motivation for entry. However, it may increase entry if only "strategic" mergers (i.e. acquisitions by large incumbents that would reinforce their market dominance) were blocked .

In this paper, we focus on the effect of digital M&As on innovation by the merging entity itself. The total innovation effect resulting from the acquisition of a start-up by a large digital platform is the combination of both positive and negative effects. Positive effects include the capacity of the acquisition to solve the "appropriability" problem of innovators who are not able to internalise all the knowledge spillovers to non-innovating firms (e.g. through imitation), which reduces their incentives to innovate in the first place (Shapiro 2011). By merging, they can internalise these externalities (Federico, Langus, and Valletti 2018). José Luis Moraga-González, Motchenkova, and Nevrekar (2022) show that the merger leads to a reallocation of the innovation effort by the merged entity among the research projects in its portfolio, which may have positive welfare effects. Next, when a merger leads to an increase in margins, the acquiring firm faces higher incentives to innovate in order to expand demand (Bourreau, Jullien, and Lefouili 2021). In addition, by merging, companies are pooling complementary skills and assets together. For instance, while the start-up might have the flexibility and reactivity to contribute innovative ideas, a large platform might be better equipped to exploit the full potential of the innovation (Crémer, Montjoye, and Schweitzer 2019, Cabral 2021).[7]

The main driver of the negative effects of M&As on innovation is their impact on the market structure. According to the the so-called Arrow replacement effect, dominant firms have intrinsically lower incentives to innovate and market power reduces innovative efforts (Aghion et al. 2005). Innovation is a competitive tool through which a firm can steal business from its competitors. By merging, previously competing firms internalise these business stealing effects, which thus reduces their incentives to innovate (Federico, Langus, and Valletti 2018; Federico, Morton, and Shapiro 2020; Motta and Tarantino 2016). A second mechanism through which M&As can deter innovation by the merging entity is the effect on the output. Innovation allows a firm to increase its margins by setting higher prices. But, in the absence of efficiency gains, M&As lead to a decrease in the merging firms' output. As a result, there is less to gain from margin-enhancing innovation (Bourreau, Jullien, and Lefouili 2021).

While a large platform may be better equipped to complete the acquired project, it may not have the incentive to develop it further (Motta and Peitz 2021, Fumagalli, Motta, and Tarantino 2020). Eventually, it may terminate the acquired project to reinforce its position on the market and be sheltered from competition. Incumbents might use acquisitions as a way to get rid of start-ups that represent potential competition because they are developing substitute products. This is documented in Cunningham, Ederer, and Ma (2021) who show that, in the pharmaceutical industry, big pharma acquires startups developing drug projects competing with their own and terminate the startup's project after acquisition, i.e. acquisition "kills" innovation.

[7] If big techs use mergers to acquire technologies, it is likely to boost the research effort of the startup (Cabral 2021) but it may reduce the organic innovation by the big tech itself. This reverse-kill phenomenon is discussed in Caffarra, Crawford, and Valletti 2020.

Several papers have tried to assess empirically the impact of mergers on innovation, by looking either at the number of patents or at the patents' citations. For instance, Haucap, Rasch, and Stiebale (2019), using data from the pharmaceutical industry, show a significant decline in the number of patents post-merger. Interestingly, the merger also negatively affects the R&D of the rivals. Fons-Rosen, Roldan-Blanco, and Schmitz (2021) compare patents belonging to acquired and non-acquired startups with similar characteristics. They find that an acquired patent's citations increase, on average, by 22% after acquisition. In their study, they compare periods of 7 years before and after acquisition but they do not look, as we do, at the evolution of citations over time. In addition, these authors did not differentiate between citations by the acquirer and citations by other firms, which we find to have different post-acquisition trends.

There are three papers closely related to ours that study the impact of mergers in digital industries. Doan and Mariuzzo (2022) analyse the cloud computing industry. They compare the innovation effort, measured by the number of patents, before and after the merger. They document an increase in the number of patents from 40% one year after the merger to 60% three years after. Accordingly, mergers seem to have a positive impact on the innovation of the merged entity, and this effect is stronger for leading firms on the market. Gugler, Szücs, and Wohak (2023) study the impact of GAFAM acquisitions on venture-capital funding and innovation, measured by patents. The main difference with our work is that they do not analyse the impact of the merger at the technology/patent level, as we do, but at a more aggregated 'market' level. These authors construct comparable groups of firms and technology classes, treated or not by the acquisition events and they estimate the impact of acquisition by comparing the two groups in a difference-in-differences set-up. They find a significant negative impact of acquisitions on venture-capital funding. The effect on innovation is less clear-cut. The initial negative effect observed for mergers before 2010 becomes positive for mergers after this date, and its magnitude varies across acquirers, with stronger effects (both positive and negative) for Microsoft. Finally, Prado and Bauer (2022) study the impact of GAFAM acquisitions on the activities of venture capital funds. They found that an acquisition by a tech giant in a given industry increases the venture capital activity in that industry with a significant increase in the number of deals and funding. However, the authors also show that this effect is only transitory and fades away after several quarters, an effect that is similar to the impact we measure on citations.

In Section 1.2, we describe the main features of Big Tech acquired technologies (1.2.1) and the construction of our working datasets (1.2.2). Sections 1.3 discusses our empirical strategy to take out the effect of endogenous factors from the technology developments around the time of acquisition. We present descriptive evidence in Section 1.4 and our main results in Section 1.5, with tests of robustness in Section 1.5.3. We develop additional analyses and extensions in Sections 1.6 and 1.7, and Section 1.8 concludes.

1.2 Empirical Methodology

In this section, we describe the data collection and the construction of the working datasets.

1.2.1 Data and Variables

We construct a database of firms acquired by Big Tech, which we match to the patents these firms filed to the US Patent and Trademarks Office (USPTO).[8] Citations to these patents can then be linked to their investigators based on the application identifiers of the citing patents.

Big Tech acquisitions

Our working sample is constructed in three steps, as presented in Table 1.1.

We first create a dataset of firms' acquisitions by Alphabet, Amazon, Apple, Meta and Microsoft. To obtain as complete of a list as possible, we merge four different databases: Standard & Poor's CapIQ (2022), Geoff, Marshall and Parker (2021), Gautier and Lamesch 2021, and the USPTO Patent Assignment Dataset (2022).[9] We retrieve

[8] USPTO-published patents represent around 82% of Big Tech-acquired patents, and 93% of Big Tech patents (as computed based on the OECD Patent Statistics, July 2021).

[9] We do not consider equity investments, licensing deals or joint ventures as acquisitions. We also do not include companies selling some of their assets as there is no transfer of the company's ownership. However, we do

information on the identities of the acquired firms and on the dates at which their acquisitions were announced. On this basis, we identify 855 public Big Tech acquisitions closed between January 1996 and January 2021 (see first column of Table 1.1).

Next, we match acquired firms with a portfolio of patents based on the name of the applicant organisation. We focus on US-granted patents,[10] which we collect from both the OECD Patent Statistics (built based on the PATSTAT database) and the USPTO Patent Views databases. By matching acquired firms with intellectual property, we can identify all (granted) patents filed by a Big Tech-acquired firm to the USPTO.

Finally, we focus on patents filed before acquisition. Patents filed under the target's name after acquisition are considered as filed by the acquirer. We find that 273 of Big Tech-acquired firms have filed at least one patent application, of which 252 before being acquired (see second and third columns of Table 1.1).

Table 1.1: Number of Big Tech-acquired firms

	Firms acquired by Big Tech btw. Jan 1996 and Jan 2021	Acquired firms with at least one US-granted patent	Acquired firms with at least one US-granted patent pre-acquisition
Amazon	105	34	27
Apple	126	53	52
Facebook	104	18	18
Google	263	75	67
Microsoft	257	93	88
TOTAL	855	273	**252**

Note: This table illustrates the steps that are taken to select, among all Big Tech-acquired firms, those that have patented a technology. Patents are identified based on their application number.

Since we identify technology developments by tracking patents as they move across firms, we will restrict our analysis to those 252 acquisitions associated with patent-protected technologies. While this only represents 29% of all Big Tech-acquired firms, this share

include companies that are only partially acquired but whose remaining assets are shut down, because the target company is no longer an independent entity after acquisition.

[10]The focus on granted patents is explained by the fact that information on the application filing date - a necessary information to derive who of the target or the acquirer filed the patent - is only available for USPTO granted patents. Because there is a lag between the filing date and the granting date, an acquired patent could be granted after acquisition (see Appendix A.1 for some details on the patenting process).

rises to 76% when we consider the biggest firms, i.e. with a total funding above $10 million (see Appendix A.2 for a graphical representation).[11]

Patent data

We collect information from Patent Views on the patents acquired by Big Tech through the acquisition of the company that filed these patents.

Patent age To control for potential trends in the technology development over a patent's life, we compute the patent age based on its filing date.

Forward citations The use and the further development of a patented technology can be proxied by forward citations received by the patent. Because 'prior art' is included in a patent by citations to previous patents, forward citations to the acquired technology reflect whether the technology is being further improved after acquisition. Appendix A.3 discusses the potential limitations attached to this use of patent citations data.

We identify the investigators of forward citations to Big Tech-acquired patents from the application identifiers of the patents containing these citations. For instance, patents cited by their acquirer can be identified by selecting all granted patents filed by Big Tech itself,[12] and then matching their application identifiers to the filing firms of the citing patents. In addition, we observe the date at which each citing patent was filed. On this basis, we can derive the number of citations received by a patent in a given month as the number of citing patents filed during that month.[13]

The most recent patents are less likely to receive citations from granted patents simply because the citing patents are not yet granted, i.e. there is a 'grant lag'. Because citations data is available until July 2022, and to avoid biases due to some citing US patents not yet being granted by that time and hence not observed, we end our study period in June

[11] Based on data retrieved from Orbis. The total funding is the sum of the variable *Shareholdersfunds* up to the year before acquisition.

[12] The Patent Views database covers all citations made by US granted patents.

[13] We assume citations are observed from the date of filing.

2017, 5 years before the data collection. Appendix A.4 illustrates that, from 2018 onwards, citing patents are indeed less likely to appear in the Patent Views database.

Acquired firms

In the end, for each patent in our database, we can identify the acquirer, the timing of acquisition, the patent's age, and the number of forward citations made every month to this patent. We construct a dataset containing all the patents belonging to Big Tech-acquired firms, and we select those firms that have published at least one patent further cited by their acquirer before July 2017. We end up with a working sample of 143 firms, i.e. 143 patent portfolios. Table 1.2 presents summary statistics on these data samples.

Table 1.2: Big Tech-acquired patents portfolios

		Portfolio size (patents #)		Patent age at acquisition (y)	
	Firms Count	Mean	SD	Mean	SD
Big Tech acquired portfolios					
AMZN	27	22.07	64.62	3.31	2.59
APPL	52	14.21	19.72	4.00	2.59
FCBK	18	7.56	17.68	4.31	4.37
GOOG	67	30.98	143.16	3.90	2.14
MSFT	88	16.52	51.88	3.86	2.79
TOTAL	252	19.84	83.12	3.87	2.71
Big Tech acquired portfolios cited by their acquirer before July 2017					
AMZN	12	15.25	24.81	3.00	2.45
APPL	29	19.79	23.49	4.83	3.22
FCBK	6	5.17	4.88	3.16	3.28
GOOG	35	56.66	195.84	5.04	2.38
MSFT	61	22.57	61.46	4.52	3.58
TOTAL	143	29.01	105.83	4.52	3.17

Notes: This table provides summary statistics on Big Tech-acquired patents portfolios.

1.2.2 Working sample

For our analysis, we construct two samples of patents. The first sample is composed of patents filed by a company later acquired by Big Tech. Our objective is to track the patented technology after its acquisition by a tech giant. We also construct a sample of comparable patents but that have not been acquired.

Acquired patents

We thus consider Big Tech-acquired patents that received at least one forward citation by their acquirer (further simply referred to as 'acquired patents'). Next, we select a balanced panel of patents observed during 4 years around the date at which the acquisition is announced.

The pre-treatment period is defined with a view to include most targets, independently from the age at which they were acquired. As such, we do not want to go back in time as far as a year before acquisition, as a significant number of (future) targets were not yet incorporated at that time,[14] and hence would not be observed over the whole study period. However, we want to be able to observe pre-acquisition potential trends in citations. With this trade-off in mind, we choose a period of 9 months before the acquisition announcement (see Figure 1.1), which allows to observe the evolution of citations before acquisition without excluding younger targets.

[14] Firms acquired within a year of their incorporation represent around 20% of Big Tech acquisitions (author's computations based on incorporation data retrieved from Crunchbase and Orbis).

Figure 1.1: Big Tech citations to acquired patents over acquisition time

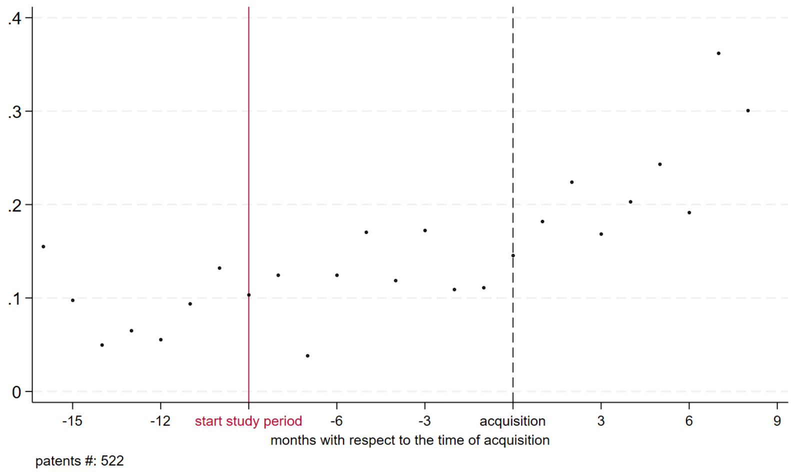

Note: The graph plots the average number of citations by the acquirer before and after acquisition.

For the post-treatment, we choose a period of 3 years after the acquisition announcement. Again, the choice of the 3 years post-treatment period is the result of a trade-off between keeping a reasonable number of observations while observing a sufficiently long period of time to analyse the dynamics of the technology developments after acquisition. Let us note that, because we end our study period in June 2017 to avoid biases in the citations count, restricting our baseline sample to patents observed up to 3 years after acquisition means that we can only use acquisitions undertaken until May 2014, which represent 58% of all 855 Big Tech acquisitions.

Of all acquired patents observed in this 4 years-window, 541 (accounting for 80 targets) are associated with at least one citation over the study period and can thus be used in our analysis of the evolution of the number of citations around acquisition.

Non-acquired patents

To control for unobserved factors that may impact the time trend in citations, we introduce a group of patents that are not treated by the acquisition event but that are comparable to Big Tech-acquired patents; namely patents that are cited by the tech giants but never acquired by them (further simply referred to as 'non-acquired patents'). These patents are assigned placebo acquisition dates by drawing from the distribution of observed Big Tech acquisitions.[15] We assume a lognormal distribution of the acquisition date $acquisition_p$ assigned to the non-acquired patent p:

$$acquisition_p \sim LN(\hat{\mu}, \hat{\sigma}^2)$$

where the mean $\hat{\mu}$ and variance $\hat{\sigma}^2$ are obtained from the distribution of observed acquisition dates.

We then select a balanced panel of non-acquired patents observed every month between 1 year since simulated acquisition and 3 years after. On this basis, we obtain two groups: i. a balanced panel of patents acquired between January 1996 and June 2017 and observed in a 4 year-window around acquisition, ii. a balanced panel of patents that were never acquired, but that have been assigned a placebo acquisition date between January 1996 and June 2017 and are observed in a 4 year-window around this placebo.

The first column of Table 1.3 presents the number of patents in these two groups: 541 patents undergo an acquisition event, and 70,136 are assigned a placebo acquisition date. The next columns of Table 1.3 present summary statistics of the citations count variable for each tech giant, separately for acquired patents and non-acquired patents. Based on a t-test at the 1% level, we find that acquired patents are on average more cited by Big Tech than non-acquired patents (with around 8 citations/acquired patent against 5 citations/non-acquired patent). This citations count variable exhibits a high variability; a majority of patents in the data set are only cited once, but a few patents are cited many times (see distribution at the monthly level in Appendix A.5).

[15] A similar study design is developed by Kleven, Landais, and Søgaard (2019), who assign placebo births to individuals who never had children by drawing from the observed distribution of age at first child among parents.

Table 1.3: Observations over the whole study period

	Cited patents	Big Tech citations				
	Count	Count	Mean	SD	Min	Max
Big Tech acquired						
AMZN	41	354	8.63	9.48	1	39
APPL	160	1,318	8.24	16.44	1	110
FCBK	7	28	4	5.51	1	16
GOOG	129	1,248	9.67	12.02	1	75
MSFT	204	1,194	5.85	16.70	1	205
TOTAL	541	4,142	7.66	15.11	1	205
Big Tech non-acquired						
AMZN	7,036	24,854	3.53	9.40	1	118
APPL	21,283	116,405	5.47	13.44	1	575
FCBK	2,455	12,946	5.27	10.69	1	105
GOOG	17,135	83,191	4.86	9.74	1	214
MSFT	29,613	99,751	3.37	9.16	1	237
TOTAL	70,136	337,147	4.81	11.60	1	598

Note: This table presents the number of observations contained in the balanced sample of patents observed in a 4 year-window around (simulated) acquisition. There are two reasons why Facebook is underepresented. First, the company is not very active from a patenting point of view. Second, Facebook started acquiring smaller firms later than the other tech giants, so most of its patented acquired technologies are not observed 3 years after acquisition.

To ensure the comparability of acquired and non-acquired patents with respect to all determinants of citations (except for the acquisition status), we use *inverse probability weighting*. This weighting consists in reinforcing the contribution of observations that are, pre-treatement, more similar to observations in the other patents group. Because most determinants of a patent's citations are unobserved, patents will be weighted directly based on the citations they received pre-acquisition. Non-acquired patents associated with the biggest weights are thus those that are, pre-acquisition, most like acquired patents with respect to their forward citations. The procedure is described in

Appendix A.6.

1.3 Model

In the previous section, we described how we collected patent citations data to capture the developments of Big Tech-acquired technologies. In this section, we make use of the time series nature of this data to identify the effect of the acquisition event.

We consider two identification strategies. First, a sharp event study, that relies on the assumption that the acquisition event is not determined by the outcome (i.e. patent citations), and on the smoothness of the average citations path absent acquisition. Second, we relax the smoothness assumption in an alternative model with a control group for acquired patents.

1.3.1 Baseline - Sharp event study

For our baseline model, we adopt a sharp event study approach as developed by Kleven, Landais, and Søgaard (2019). The development of the acquired technology by the acquiring firm is measured by citations to the associated patents. We study the evolution of the number of forward citations by the acquirer as a function of event time dummies, which represent the quarters (three months) in which citing patents are filed with respect to the time of acquisition $t = 0$.[16] To identify the impact of a Big Tech acquisition, we must correct for the potential endogeneity coming from determinants of the technology development other than acquisition. Most of these determinants are unobserved or even unknown, but we could indirectly control for them by introducing life-cycle trends (e.g. the number of forward citations might depend on the stage of a patent's life) and business-cycle trends (e.g. the industry's R&D might be more or less dynamic in given years).

We denote by $Cit_{p,j,t,d}$ the number of forward citations to patent p of the target

[16] The event time dummies are constructed by situating the month in which the patent is filed with respect to the month in which it is acquired and, to limit variability, aggregating by quarter: $t \in \{-3 = (-10m, -9m, -8m), ..., 0 = (-1m, 0m, 1m), ...\}$ with $0m$ when the filing month coincides with acquisition.

firm j at event time t and date d. Target-specific fixed effects are captured by $firm_j$. We control for life-cycle trends and business-cycle trends by including the patent's age $age_{p,d}$ and a full set of calendar date d dummies in the vector M' ($d = 1996q1, 1996q2, \dots, 2017q2$).[17] The effects of all included regressors are identified because patents are acquired at different times; conditional on date and age, there are variations in event time. We define the following model:

$$Cit_{p,j,t,d} = f(J'\theta^1, firm_j\xi^1, age_{p,d}\beta^1, M'\gamma^1) + \varepsilon^1_{p,j,t,d} \tag{1.1}$$

where J' is a vector containing the time dummies at the quarterly level ($t = -3, \dots, -1, 0, 1, \dots, 12$) excluding the base category $t = 0$.

To define the function $f(.)$, we must account for the nature and distribution of the response variable: the citations count. The most widely used model for a count regression is the Poisson distribution. However, the Poisson model assumes that the mean and variance of the errors are equal. In our case, the variance of the citations count is much larger than its mean: a majority of patents in the data set are only cited once, but a few patents are cited many times. Fitting a negative binomial model is a way to correct for the over-dispersion observed in the distribution of the citations count variable (Ajiferuke and Famoye 2015).[18]

The negative binomial distribution function of the citations count can be written as:

$$P\left(Cit = Cit_{p,j,t,d} \mid t, age_{p,d}, d, firm_j\right)) =$$
$$\begin{pmatrix} 1/\delta + Cit_{p,j,t,d} - 1 \\ Cit_{p,j,t,d} \end{pmatrix} \left(\frac{\delta\mu\left(t, age_{p,d}, d, firm_j\right)}{1 + \delta\mu\left(t, age_{p,d}, d, firm_j\right)}\right)^{Cit_{p,j,t,d}} \left(\frac{1}{1 + \delta\mu\left(t, age_{p,d}, d, firm_j\right)}\right)^{1/\delta}$$

where $\mu(.)$ is the mean of the model and δ is the dispersion parameter, which accounts for a variance of the data that is higher than the mean, and $Cit_{p,j,t,d} = 0, 1, 2, \dots$.

[17] The calendar date dummy is defined as the quarter associated with the month in which the citing patent is filed, e.g. $(2013m7, 2013m8, 2013m9) = 2013q3$.

[18] We test whether the Negative Binomial model is appropriate by comparing it to a Poisson model using the likelihood ratio test. We find that the δ dispersion parameter for model 1.1 is significantly different from zero (χ^2 = 2985), which contradicts the assumption of the Poisson model. On this basis, we can confirm that a Negative Binomial regression should be used.

On this basis, we identify the changes in the acquired technology development that can be attributed to a Big Tech acquisition as the changes in citations with respect to the time of acquisition. Because the negative binomial model is used, $\hat{\theta}^1_t$ identifies the expected difference in log citations between quarter t and the reference group ($t = 0$): $\hat{\theta}^1_t = ln(Cit_{p,j,t,d} \mid age_{p,d}, d, firm_j) - ln(Cit_{p,j,0,d} \mid age_{p,d}, d, firm_j)$. To obtain a more intuitive interpretation of our results, we will use the incident rate ratios: $e^{\hat{\theta}^1_t} = \frac{Cit_{p,j,t,d} | age_{p,d}, d, firm_j}{Cit_{p,j,0,d} | age_{p,d}, d, firm_j}$. By taking the exponential function, the difference in log citations becomes the ratio of the citations count at a given event time to the citations count at acquisition. The validity of the approach is further discussed in Appendix A.7.

1.3.2 Introducing a control group

While life-cycle and business-cycle trends can be directly controlled for, some other determinants of the technology development are unobserved (e.g. upward trends in forward citations due to technology spillovers). To disentangle the cross-sectional correlation in the data from the effect of acquisition, we introduce a control group not treated by the acquisition event: Big Tech-cited (but never acquired) patents. These patents are assigned placebo acquisition dates randomly drawn from the distribution of observed acquisitions by assuming a lognormal distribution (as described in Section 1.2.2). We rewrite model 1.1 as follows:

$$Cit_{p,t,d} = f(J'\theta^2, A_p\iota^1, J'A_p\alpha^1, age_{p,d}\beta^2, M'\gamma^2) + \varepsilon^2_{p,t,d} \tag{1.2}$$

where $A_p = 1$ if patent p is acquired, $A_p = 0$ otherwise.[19]

On this basis, we can estimate the impact of Big Tech (simulated) acquisition for both acquired and non-acquired patents separately. If life-cycle and business-cycle trends captured all determinants of the evolution of citations other than acquisition, the impact of acquisition for non-acquired patents after controlling for age and date should be null. In other words, the trend in citations to non-acquired patents over event time captures

[19] In this second model specification, firms fixed effects are no longer accounted for as we cannot retrieve the identities of all the firms cited by Big Tech patents.

the remaining unobserved heterogeneity. The effect of acquisition can therefore be estimated as the event time impact for acquired patents with respect to non-acquired patents. When the outcome variable is negative binomial-distributed, this can be estimated by the "Difference-in-semielasticities" (DIS),[20] i.e. the acquisition status' impact on the semielasticity of citations with respect to the event time: $e^{(\widehat{\theta}_t^2+\widehat{\alpha}_t^1)} - e^{(\widehat{\theta}_t^2)}$. The validity of the identification parallel trends assumption can be verified from the pre-acquisition DIS.

1.4 Preliminary analysis: Comparing average citations counts

To start with, we present some preliminary evidence on the evolution of citations after acquisition.

A patent receives, on average, 0.09 citation/month before being acquired and 0.18 citation/month after. This increase in citations after acquisition suggests that the acquiring firm invests in the technology of the acquired firm and continues to develop it after acquisition.

Citations thus appear on average twice as high after acquisition than before. To make the same comparison absent life-cycle and business-cycle trends, we define simplified versions of models 1.1 and 1.2, with the dummy variable *Post* taking the value 1 after acquisition:

$$Cit_{p,j,d} = f(Post\ \theta^1, firm_j \xi^1, age_{p,d}\beta^1, M'\gamma^1) + \varepsilon_{p,j,d} \tag{1.3}$$

$$Cit_{p,d} = f(Post\ \theta^2, A_p \iota^1, Post\ A_p \alpha^1, age_{p,d}\beta^2, M'\gamma^2) + \varepsilon_{p,d} \tag{1.4}$$

The estimation results are presented in Table 1.4. The parameter estimates based on the sample of acquired patents only (Model 1.3) show a significant increase in citations after acquisition. The model estimates that an acquired patent receives 35% (IRR = $e^{\hat{\theta}^1}$ =

[20] When the conditional mean function is non-linear, the parameter associated with the interaction term does not provide a consistent estimate of the interaction effect (Shang, Nesson, and Fan 2018).

1.35*** (0.10)) more citations by its acquirer after the acquisition. The results for Model 1.4 on (unweighted) acquired and non-acquired patents are similar, with an estimated citations increase of 48% (DIS = $e^{(\widehat{\theta^2}+\widehat{\alpha^1})} - e^{(\widehat{\theta^2})}$ = 0.48*** (0.12)) after acquisition.

Table 1.4: Big Tech citations to acquired and non-acquired patents

	Model (3)	Model (4)
Post	.30***	.15***
	(.07)	(.01)
Acquired		.50***
		(.06)
Post#Acquired		.34***
		(.07)
Firms FE	Yes	No
Date dummies and Age	Yes	Yes
Patents #		
acquired	541	541
non-acquired		77,522

Std. err. in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

This preliminary evidence tends to show that the target technology development by the acquirer increases significantly after acquisition. In other words, that the acquirer is doing significantly more research effort to develop the acquired technology. In the next section, we refine this analysis by reinforcing the contribution of non-acquired patents that are more similar to acquired patents, and by allowing the effect of acquisition to vary over event time.

1.5 Impact of acquisition on the acquired technology

In this section, we present our main results regarding the impact of a Big Tech acquisition on the development of the acquired technology, as measured by citations to the associated patents. Model 1.1 is estimated on the balanced panel of Big Tech-acquired patents. Model 1.2 is estimated on the balanced panel of trimmed Big Tech-acquired and non-acquired patents weighted based on their inverse probabilities. These models allow us to track the evolution of citations over time and give a more accurate view of the technology development by the acquirer after acquisition.

1.5.1 Results: Baseline - Sharp event study

We estimate our models by including the full set of time dummies (at the quarter level). This allows us to see the evolution of citations up to three years after acquisition. The results of Model 1.1 are presented on Figure 1.2. On the figure, we represent the estimated incident rate ratios ($e^{\hat{\theta}^1_t}$) for each quarter and we include 95% confidence bands around the event coefficients. We control for life-cycle and business-cycle trends and for the acquired firm fixed effect. The estimated coefficients represent the ratios of the citations count for each event time to the citations count at acquisition. A value above 1 means that citations increase after acquisition.

Our results confirm the preliminary evidences that acquisition increases citations but now we can identify that this increase is only *temporary*. Citations experience a non-lasting boom after acquisition. Looking at the results in more details, on Figure 1.2, we observe that citations increase significantly up to 1.5 year after acquisition (citations then appear to be more than 50% higher compared to their acquisition level). After that, citations start slowing down. The evolution of citations by the acquirer thus follows a bell curve and, as we will show, this result is robust to many alternative specifications. These results suggest a continuous development of acquired technologies but the R&D effort of the acquirer is fading away after some time.

Figure 1.2: Big Tech citations to acquired patents relative to acquisition

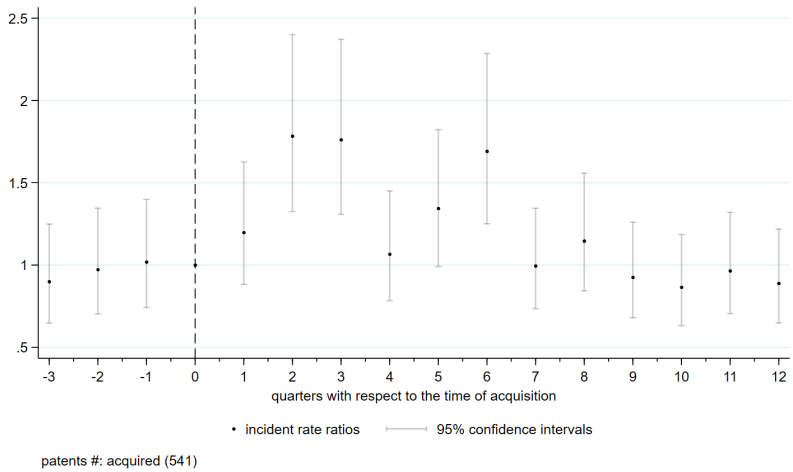

Notes: The graph shows the incident rate ratios for acquired patents: $e^{\hat{\theta}_t^1}$ from model 1.1. These coefficients are estimated on a balanced sample of patents in a 4 year-window around acquisition.

Since the impact of acquisition is identified from the sharp breaks in citations trajectories immediately following acquisition, our identification strategy can handle the smooth trend in citations which, even if not significant, appears slightly positive in the quarters before acquisition. In the next section, we propose an alternative identification strategy, with which we aim to take out the citations trend (even smooth) coming from factors exogenous to the acquisition event.

1.5.2 Results: Introducing a control group

We present on Figure 1.3 the DIS estimated based on Model 1.2. For this estimation, we use the balanced panel of trimmed Big Tech-acquired and non-acquired patents. The contribution of each observation has been multiplied by its inverse probability weight. These estimates can be interpreted as the changes in the number of acquirer's citations at event time t relative to the acquisition time, having controlled for life-cycle and business-cycle trends, for acquired patents with respect to non-acquired patents. A value above 0 means that, compared to non-acquired patents, citations to acquired patents are higher than at acquisition.

In support of the assumption that citations to acquired and non-acquired patents (conditional on the propensity scores) would move in parallel absent acquisition, we see that the DIS are insignificant in the pre-acquisition period. Just after acquisition, we see that Big Tech citations grow faster for acquired patents than for non-acquired patents (independently from life-cycle and business-cycle trends), identifying a boost in the development of acquired technologies by the acquiring platform. After 1.5 year, these technology developments start slowing down, suggesting that the boost in the acquired technology development fades away in the long run.

Figure 1.3: Big Tech citations to acquired patents w.r.t. non-acquired patents, relative to the (simulated) acquisition announcement

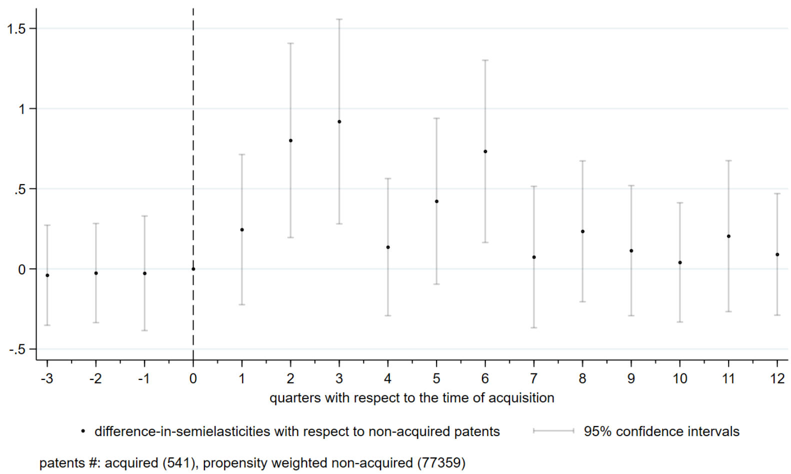

Notes: The graph shows the DIS between acquired and non-acquired patents: $e^{(\widehat{\theta}_t^2+\widehat{\alpha}_t^1)} - e^{(\widehat{\theta}_t^2)}$ from model 1.2. These coefficients are estimated on a balanced sample of patents in a 4 year-window around (simulated) acquisition.

The results of our different models are convergent and they show that citations experience a boom after acquisition. We interpret this as an increased research effort by the acquirer to further develop the technologies it acquires. However, this boom in the acquirer's R&D activity is not lasting and, after 1.5 year, the identified effect fades away. In the next section, we will show that this inverse U-shaped trend is robust to alternative model specifications.

1.5.3 Robustness checks

To test for the robustness of our results, we replicate our baseline analysis with alternative model specifications. These robustness checks are presented in Appendix A.8.

Additional regressors First, we propose to replicate our analysis with more citations determinants included as regressors in the model. In particular, we control for the acquirer's identity (Microsoft versus others), for the technology field to which the acquired patent belongs and for the origin of the publishing company.

Alternative study periods Second, we replicate our analysis based on alternative study periods. First, we change the study period by extending the pre-treatment period from 3 to 5 quarters (15 months before acquisition). Second, we reduce our study period to 2 (instead of 4) years around acquisition. Third, we reduce our post-treatment period such as to include *Motorola Mobility*, which was acquired by Google but later sold to Lenovo and hence not included in our baseline sample.

Heterogeneity-robust treatment effects Finally, we aim to obtain estimates that are robust to heterogeneous treatment effects (see Borusyak, Jaravel, and Spiess 2024; De Chaisemartin and d'Haultfoeuille 2024; Sun and Abraham 2021), i.e. to potential variations in the effect of acquisition across acquisition dates. To obtain heterogeneity-robust estimates in our non-linear context, we adopt the approach developed by Wooldridge (2023). We first group patent portfolios based on the date in which they were acquired to construct 'acquisition-cohorts'. We then estimate the treatment effects for each cohort separately, and we aggregate all these cohort-specific treatment effects to obtain the average treatment effects.

In all these alternative specifications, we find results that are consistent with our baseline estimates, with a significant but non-lasting boost in citations after acquisition.

1.6 Technology development by non-acquiring firms

To explain the observed slow down in citations after acquisition, we put forward the hypothesis of technology maturity; the acquired technology could be less developed because of diminishing returns to the innovative effort. According to this hypothesis, the tech giants are acquiring technologies that are close to maturity. By pooling skills and assets following acquisition, they manage to complete the development of the technology, which is not further developed but, instead, directly included in a product. In other words, the development slows down because the technology reaches its maturity.

A corollary of this hypothesis is testable: if the slow down of the acquired technology development by its acquirer was explained by technology maturity, we should observe a similar slow down of its development by the rest of the industry. In our patents context, this would translate into citations by the acquirer and citations by the rest of the industry following a similar pattern.

We estimate model 1.1 on two separate samples: Big Tech-acquired patents cited by their acquirer, and Big Tech-acquired patents cited by other firms than their acquirer. Out of the 541 Big Tech-acquired patents in our sample, 484 are also cited at least once over our study period by other firms than their acquirer. The estimated incident rate ratios ($e^{\hat{\theta}_t^1}$) are presented on Figure 1.4, separately for these two citing groups.

On Figure 1.4 (b), we observe that the acquisition by a tech giant induces a positive effect on the citations by non-acquiring firms; they increase by up to 50% after acquisition. The acquisition acts as a signal, putting the acquired firm in the spotlight and boosting the research effort in its technology field. However, while citations by the acquirer already start slowing down after 1.5 year, citations by other firms than the acquirer keep increasing up to 2.5 years after acquisition.

Figure 1.4: Citations to Big Tech-acquired patents relative to acquisition

(a) by the acquirer (b) by other firms

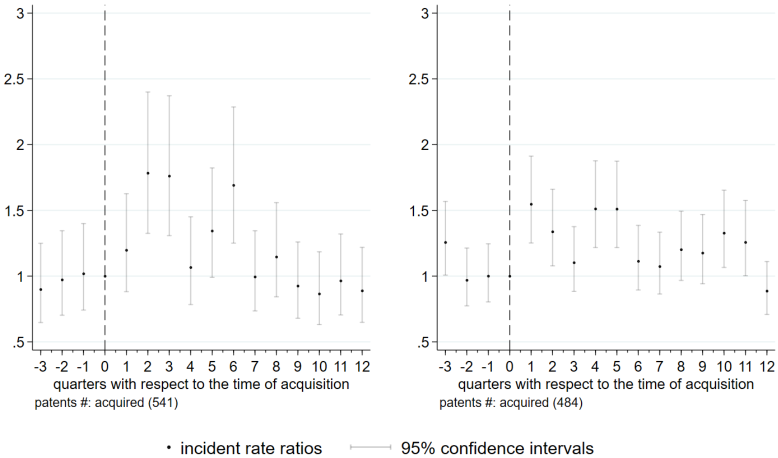

Notes: The graph shows the incident rate ratios for acquired patents: $e^{\hat{\theta}^1_t}$ from model 1.1. These coefficients are estimated on a balanced sample of patents in a 4 year-window around acquisition.

These results do not corroborate the technology maturity hypothesis. The rest of the industry continues to invest to develop the acquired technology while the acquirer's innovative effort has already started fading away.[21] This suggests that, at the time we observe a decline in the acquirer's effort to develop the acquired technology, this technology's improvement potential has not been fully exhausted. So technology maturity alone does not seem to provide a credible explanation for the slow down of the acquired technology development by its acquirer.

[21]This positive impact of acquisition on the non-merging parties is consistent with the model of Federico, Langus, and Valletti (2018). As a response to acquisition, the rest of the industry does more research effort, possibly to catch-up and to compensate for the disappearance of the independent startup. Let us also note that this contrasts with Affeldt and Kesler (2021)'s finding that outsiders invest less in the product - in their context, an app - development after its acquisition by a tech giant.

1.7 Extension: Effects across portfolio sizes

In this section, we extend our baseline results on the technology development by its acquirer. We aim to investigate whether the effect of acquisition varies with the size of the acquired patents portfolio. To do that, we refine model 1.1 by allowing the event time impact to vary with the size of the target's patents portfolio:

$$Cit_{p,j,t,d} = f(J'\theta^5, Large_p\zeta^2, J'Large_p\eta^2, age_{p,d}\beta^5, M'\gamma^5, firm_j\xi^3) + \varepsilon^3_{p,j,t,d} \quad (1.5)$$

where $Large_p$ takes the value 1 if patent p belongs to a large portfolio.

In our sample, almost half of the observations belong to a portfolio with 32 or more published patents. We therefore identify a large acquired portfolio as containing at least 32 patents. We use a second measure based on a cutoff value of 5 patents for the portfolio size. In this second specification, patent p belongs to a large portfolio if it contains more than 5 patents, with a majority of patents falling in this category.

The estimated incident rate ratios are presented on Figure 1.5. We see that technologies belonging to small portfolios are more developed by their acquirer than technologies belonging to large portfolios. For patents in a portfolio with 32 or more patents, the boost in the acquirer's citations just after acquisition is insignificant, while a negative effect is observed after 1.5 year. On the contrary, for patents in smaller portfolios, the decline in citations is less pronounced and the effect remains positive even at the end of the study period. The alternative definition of a large portfolio gives similar - although less pronounced - results.

An intuitive interpretation of this result is that, for small targets, the acquisition of a specific technology explains a significant share of the acquisition decision while, for large targets, a bigger share of the acquisition decision is left unexplained, i.e. many patents in a large portfolio may be of little interest to the acquirer.[22] This suggests that the acquisition of small portfolios is more likely to have been driven by a specific patent,

[22] Let us however remind the reader that patents should be cited at least once by the acquirer to be included in our sample.

which the acquirer invests significant effort into further developing.

Figure 1.5: Big Tech citations to acquired patents relative to acquisition, by target's size

(a) <= 32 vs > 32 patents

(b) <= 5 vs > 5 patents

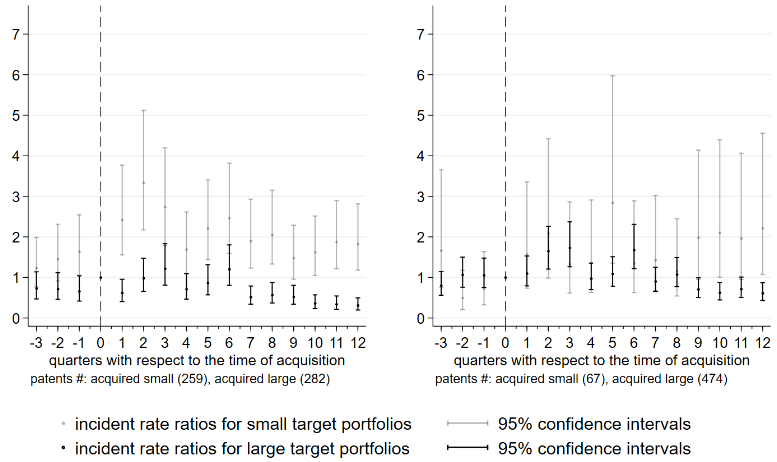

Notes: The graphs show the incident rate ratios from model 1.5 for small (e^{θ^5}) and large ($e^{\theta^5+\eta^2}$) acquired portfolios.

1.8 Conclusion

With this paper, we aim to bring empirical evidence of the effect of acquisitions in digital markets on acquired innovative technologies. Information provided by the patent system allows us to track technologies before and after they are bought by 'Big Tech', i.e. Alphabet, Apple, Meta, Amazon and Microsoft. To study the development of an acquired technology, we use information on citations made to the patents protecting that technology in subsequent patents. Accordingly, the development of Big Tech-acquired technologies by their acquirer is proxied by Big Tech's citations to acquired patents.

Just after acquisition, we find a positive effect of acquisition on the improvements made by Big Tech to acquired technologies. After 1.5 year, these developments of the acquired technology by the acquiring platform start slowing down. A potential explanation for this result is that the acquired technology reaches full maturity thanks to the pooling of skills and assets of the digital platform and the acquired start-up. However, these Big-Tech acquired technologies keep being improved by the rest of the industry up to 2.5 years after acquisition, which means that their improvement potential has not been fully exhausted. On this basis, we conclude that technology maturity alone cannot explain the slow down in the development of Big Tech-acquired technologies.

More generally, our analysis contributes to the understanding of the impacts of mergers and acquisitions on the evolution of the acquired products and technologies, a research field where empirical evidence remains scarce. We have chosen to focus our analysis on acquisitions by Big Tech, mainly because of the very high rate at which these platforms have acquired start-ups in the past twenty years. Our conclusions are thus based on acquisitions by dominant firms, mainly in the digital sector. Future work could have a larger focus, including less powerful acquirers and more diverse industries.

Chapter 2

Talent Acquisition and Technology Leadership: A Study of Digital Platforms[1]

Abstract

In this paper, we track inventors from firms acquired by Big Tech platforms, and we identify those who continue to innovate for their acquirer post-acquisition. Drawing on firm-level and patents databases, we find that inventors working pre-acquisition on technologies that are closer to their acquirer's "core" fields are more likely to keep innovating for their acquirer. We extend our analysis by investigating hypothetical explanations for this finding. First, we find that inventors specialising in core technologies are easier for the acquirer to assess, which could partially explain that they are more likely to stay active after acquisition. However, post-acquisition, these talented core inventors are often directed away from their acquirer's core technology fields. These results suggest that Big Tech acquires inventors in its core technology fields mainly because it can better evaluate their value, rather than to keep them focused on those core technologies.

2.1 Introduction

In the digital economy, knowledge is what creates value. As an alternative to internal R&D efforts, companies can generate this knowledge by leveraging the expertise of other

[1]This chapter is co-authored with Axel Gautier.

firms. When it is patented, knowledge can be accessed through licensing agreements or acquired via patent assignments. It can also be copied (e.g. reverse engineering, industrial espionage). But these are temporary measures, as patents expire and technologies evolve.[2] The talent of the people who generated the knowledge, on the other hand, is an asset that can produce a continuous flow of future knowledge. In this paper, we are interested in how firms integrate such talent into their own research teams to access external knowledge.

Directly hiring talent from a competitor is not always easy. Within Silicon Valley, entrepreneurs might fear for their reputation and investor relationships (Polsky and Coyle 2013). Knowledge is also sometimes embedded in a team of people and could dissipate if the team is torn apart (Jaravel, Petkova, and Bell 2018). Acquiring the entire firm instead could be more efficient in accessing its talent. The practice of acquiring a firm with the main aim to access its know-how and human capital is often referred to as 'acqui-hiring' (Polsky and Coyle 2013).[3]

An acquirer who successfully retains the acqui-hired staff and integrates them effectively can benefit from knowledge transfer and access to expertise.[4] The practice allows to quickly onboard an entire team through one transaction (J. D. Kim 2020), and to avert knowledge leaks that could accelerate the depreciation of the acquired human capital (Selby and Mayer 2013). Entrepreneurs themselves may also prefer acqui-hiring over defection because selling a startup to a leading company often elevates their status within Silicon Valley (Polsky and Coyle 2013), and because it may be financially advantageous compared to direct talent recruitment due to the favorable tax treatment on capital gains versus compensation income (Polsky and Coyle 2013).[5]

[2]Higgins and Rodriguez (2006) find that companies facing expiring patents are more likely to engage in mergers and acquisitions.

[3]In an 'acqui-hiring' context, the acquiring firm is primarily seeking experience, expertise and engineering talent, not necessarily an existing product (Varian 2021).

[4]In R&D-intensive industries, the retention of employees' tacit knowledge is key (Grant 1996). But there are many reasons why the acquirer might not be able to retain or properly integrate the acqui-hired staff, see for instance Block, Fisch, and Van Praag (2017); Campbell et al. (2012); Michaels, Handfield-Jones, and Axelrod (2001); Ranft and Lord (2000); Ranft and Lord (2003); Seitz and Kopp (2023) and Xiao and Dahlstrand (2023).

[5]Despite these advantages for the parties to the merger, scholars warn of the inefficiencies associated with the concentration of inventors in large incumbents, where their innovative output is expected to be lower than

For Alphabet, Apple, Meta, Amazon and Microsoft (aka "Big Tech"), gaining access to the target's talent is an important driver of start-ups acquisitions. Acquiring the talent from innovative start-ups is a way for incumbents to source knowledge that is critical to maintain an existing technological leadership[6] (Puranam, Singh, and Zollo 2006) or, instead, to venture into new product areas where the acquirer's own experience is limited (Varian 2021). With this paper, we aim to provide an empirical analysis of the strategies behind the practice of talent acquisition along these goals of technology leadership and diversification. To do so, we want to go further than testing whether or not the target inventor keeps working at the acquirer's (like in studies focusing on the retention rate[7]). Instead, we want to study the profiles of inventors who keep *innovating* for their acquirer after acquisition. We ask: Are inventors of technologies in which Big Tech is already strong more likely to further innovate for their acquirer?

To answer this question, we first construct a database of all the firms acquired by Big Tech, which we match to firm-level data from Orbis Global and to patent data from the USPTO PatentsView database. We then develop our analysis at the level of inventors from Big Tech-acquired firms. We identify these inventors from the patents they filed at Big Tech-acquired firms, so we focus on inventors who have filed at least one patent application before being acquired. Among these target inventors, we want to identify those who will develop intellectual property for their acquirer. To do so, we check when the name of an inventor previously filing patents for the target is found after acquisition in a patent filed by the acquirer. And we can test whether this varies with the nature of the innovation developed by the target inventor pre-acquisition. We find that 74% of Big Tech targets are associated with at least one inventor further patenting for their

in young firms (Akcigit and Goldschlag 2023), and with "talent hoarding", when acquisition aims to deprive a competitor in which the startup employees would have been more productive (Benkert, Letina, and Liu 2023; Bryan and Hovenkamp 2020a) and obtain monopsony power over specialised talent (Bar-Isaac, Johnson, and Nocke 2023).

[6]Such practice could be viewed as a substitute to the - much less desirable - exclusionary practices (e.g. reducing interoperability with the startup's product or by imitating its main features) in which digital platforms engage to drive startups away from their core market (Motta and Peitz 2021, Shelegia and Motta 2021).

[7]See for instance Seitz and Kopp (2023), who hand-collected data on founders' retention, Xiao and Dahlstrand (2023), who used employer-employee data of the Swedish high-tech sectors, and Ranft and Lord 2003, who ran a survey of managers to identify the retention of key employees.

acquirer, and that inventors working pre-acquisition on technologies that are closer to their acquirer's "core technologies" (i.e. patented technologies of which a tech giant owns a significant share) are more likely to keep innovating for their acquirer after acquisition.

Next, we test two hypotheses that could explain this result. We first suggest that Big Tech might more easily assess the value of inventors when they work on its core technologies, and hence better exploit their talent, while talented inventors active in technologies outside of their acquirer's core business might not be recognised for their true worth. This hypothesis is supported by the finding that, for the latter "non-core" inventors, talent does not translate into a higher probability to keep innovating for their acquirer. However, once acquired, more talented core inventors seem to move away from their acquirer's existing core businesses. These results suggest that Big Tech acquires inventors in its core technology fields mainly because it can better assess their value, rather than to maintain their focus on those core technologies.

Related literature

Our paper is related to the literature on the interactions between the technologies of merging firms.

At the product level, theory predicts that the merger between substitute technologies should lead lower ex-post R&D efforts due to the termination of concurrent R&D programs (Cassiman, Colombo, et al. 2005), while merging complementary technologies brings about synergies and economies of scope, which should in turn stimulate the R&D process (Caves 1989; Cohen and Levin 1989; Röller, Stennek, and Verboven 2006). At the firm level, empirical studies of high-tech industries usually find that the greatest innovative output occurs when there is a high degree of overlap in the technological profiles between the acquirer and the target (Cloodt, Hagedoorn, and Van Kranenburg 2006; Kapoor and Lim 2007). However, in the case of Big Tech acquisitions, the acquirer is disproportionately bigger than its target so the effect on the aggregated R&D output is unlikely to be observable. What can be observed, instead, is what happens to the acquired knowledge, and how this relates to the ex-ante technological similarities between the acquirer and the target.

Mergers and acquisitions are a crucial means for both enhancing and diversifying a firm's internal knowledge base (Rosenkopf and Almeida 2003, Zheng, Ulrich, and Sendra-García 2021). The acquisition of related firms complements the acquirer's internal knowledge base, whereas acquiring unrelated firms broadens the scope of knowledge (Lodh and Battaggion 2015). Case studies of Big Tech acquisitions usually find that acquirers are strengthening their own segments. For instance, Argentesi et al. (2021) find that acquisitions by Meta and Amazon focus on areas of economic activity in which they are already strong. Similarly, Gautier and Lamesch (2021) find that Microsoft, Amazon and Apple mostly acquire products targeted to their main user segment. This might be explained by the fact that companies with a technologically similar patent portfolio can better integrate the target's knowledge (Ahuja and Katila 2001; Arts, Cassiman, and Hou 2021; Cassiman and Veugelers 2006). Similarly, technological similarities between the acquirer and the target are expected to play a role in the future of acquired inventors.

On the one hand, the acquisition of smaller firms brings in specialized expertise and talent that can help the acquirer expand to areas where it has limited experience (Varian 2021). On the other hand, the acquirer might better assess and integrate new ideas when the acquired inventors' expertise aligns with its existing R&D portfolio. In an empirical study in the field of biotechnology, Verginer et al. (2022) find that inventors who operate within the same patent classes as their acquirer tend to maintain higher productivity levels following the acquisition. In addition, from the inventor's perspective, an expertise distinct from that of the acquiring firm could be more valued outside of the merged entity and hence increase the likelihood of departure post-acquisition (Verginer et al. 2022). With this paper, we want to test how the findings of this literature extend to the digital world, and whether Big Tech's target inventors whose expertise aligns with their acquirer's existing R&D portfolio are more or less likely to keep innovating for their acquirer.

We describe the construction of our dataset and the variables of interest at the firm, inventor and technology levels in Section 2.2. Section 2.3.1 sets up our research questions and empirical strategies, of which Sections 2.4.1 and 2.4.2 discuss some extensions. Section 2.5 concludes.

2.2 Data

2.2.1 Firms' characteristics

We identify firms acquired by Big Tech and the dates at which their acquisitions were announced from four different databases: Standard & Poor's CapIQ (2022), Geoff, Marshall and Parker (2021), Gautier and Lamesch (2021), and the US Patent and Trademarks Office (USPTO) Patent Assignment Dataset (2022). On this basis, we identify 859 firms acquired by Big Tech between January 1996 and January 2021.

Firm level data is collected from the Orbis Global database. Out of the 859 Big Tech-acquired firms that we identified, 395 can be found in Orbis Global.[8] We collect information on the firm's country of origin, number of employees, incorporation date, and funding amount. Orbis contains some missing values for the number of employees and the funding amount; we add an online search step to try and fill them in. Table 2.1 presents some summary statistics on the obtained dataset.

The distribution of the number of employees and of the funding data is highly skewed (see Appendix B.1). To better capture the information contained in the number of employees data, we define a variable $FirmSize^T$ based on the OECD thresholds: small firms with < 50 employees, medium firms with $[50; 250[$ employees, and large firms with $>= 250$ employees. Similarly, we split the sample into three categories based on the quantiles of the funding data: $Funding^T \in \{\leq \$230M;\ ]\$230M, \$1,730M];\ > \$1,730M\}$.

[8] Targets that are located in Asia, Canada or South America are less likely to be found in Orbis, as well as targets acquired before 2006.

Table 2.1: Big Tech acquired firms

	Origin		EmplNbr			Funding ($1M)			Incorp	
	obs.	US = 1	obs.	Mean	SD	obs.	Mean	SD	obs.	[min,max]
AMZN	45	76%	24	3877	17716	24	1570	7306	47	1978, 2020
APPL	59	58%	38	708	86	49	56	130	61	1992, 2019
FCBK	50	62%	28	194	797	28	24	45	53	1989, 2020
GOOG	107	76%	59	422	2472	56	143	596	112	1967, 2020
MSFT	121	73%	60	293	1237	65	259	1516	122	1978, 2019
TOTAL	382	71%	209	687	6186	222	297	2553	395	1967, 2020

This table presents some summary statistics on Big Tech targets matched with Orbis data. $US = 1$ when the Orbis variable *CountryISOcode* is associated with the value "US", $\boldsymbol{EmplNbr}$ is the value taken by the Orbis variable *Numberofemployees* in the year before acquisition, and $\boldsymbol{Funding}$ is the sum of the Orbis variable *Shareholdersfunds* up to the year before acquisition.

2.2.2 Inventors' characteristics

To collect information on inventors acquired by Big Tech, we will be using the content of patent documents, with a focus on US patents filed by Big Tech platforms (i.e. the acquirers) and by Big Tech-acquired firms (i.e. the targets). Patents are attributed to the acquirer if they are filed under the acquirer's name, or if they are filed under the target's name after acquisition.[9]

Patent data is collected from the USPTO PatentsView database (e.g. patent number, patent date, application identifier, publication author, CPC technology field and inventor's names[10]). Out of the 859 identified Big Tech's targets, 252 have filed at least one patent application before being acquired.

Patent information

The USPTO PatentsView database contains patent-level data (see Appendix B.2 for descriptive statistics) that can be used to describe the profiles of target inventors. We con-

[9] Around 5% of the patents associated with a target's name are filed after acquisition. In a robustness check, we will further only attribute to the acquirer those patents filed under the acquirer's name.

[10] Harmonized inventors' names can be found at https://patentsview.org/download/data-download-dictionary.

struct the following variables at the inventor-level:

- the total number of patents in which the inventor is listed,
- the number of patents in which the inventor is listed as first author,
- the number of patents originating from the US in which the inventor is listed,
- the number of patents in which the inventor is listed together with some co-author,
- the number of months since the inventor's first filing before acquisition,
- the number of months since the inventor's last filing before acquisition.

These variables are constructed over all the patents filed by a given inventor for the target firm.

Technology leadership

We aim to determine whether the acquired inventor was working on a technology in which her acquirer is focusing its own innovative efforts. To do so, we identify technologies belonging to her acquirer's "core" fields in the year of acquisition.

Core technologies are patented technologies of which a tech giant owns a significant share. This is translated in our patents context as: a significant share of the total number of patents classified in that technology field. So core technologies are defined based on the relative frequency at which they appear in a tech giant's patent portfolio.

The technology fields to which a patent belongs is recorded in the CPC classification, which contains 131 subsections at the two-digits level.[11] To identify Big Tech "core" fields, we first compute, at the yearly level, the number of patents classified in each of the 131 cpc subsections across the whole USPTO patent database. Next, we compute the share represented by each tech giant's portfolio in these respective subsections.[12] We consider that a given cpc subsection represents a Big Tech core technology field, in a

[11] See detailed list: https://www.uspto.gov/web/patents/classification/cpc/html/cpc.html

[12] Most Big Tech's patents are classified in the Physics Section, e.g. 'Computing; Calculating or Counting' (CPC G06).

given year, if the Big Tech's portfolio contains at least 1%[13] of all occurrences of that technology in that year. Because big tech platforms might change the fields in which they specialize over time, this definition allows a dynamic setting in which we can identify whether the acquired technology belongs to a core field at the time of acquisition.

Based on this information at the patent-level, we can construct a variable $\overline{Core_i}$ at the inventor-level, defined as the share of core fields averaged across all the patents filed by inventor i:

$$\overline{Core_i} = \frac{1}{P}\sum_p \frac{1}{F}\sum_f I_{f,p,i}$$

where $I_{f,p,i} = 1$ if a given field f listed in the patent p filed by the inventor i is associated with at least one of the acquirer's core fields in the year of acquisition.

On average, 34% of the fields listed in patents filed by a target inventor belong to one of her acquirer's core technologies at the time of acquisition. However, we observe very high peaks on the upper and lower bounds of the distribution of the $\overline{Core_i}$ variable: out of 4,966 target inventors,[14] 1,077 (22%) are associated exclusively with core technologies, and 2,573 (52%) are not associated with any core technology.

For this reason, we also construct a binary variable $Core_i$, capturing whether an inventor i is active in her acquirer's core technology fields:

$$Core_i = max\ I_{p,i}$$

where $I_{p,i} = 1$ if a given patent p filed by the inventor i is associated with at least one of the acquirer's core fields in the year this patent was acquired.[15]

[13] Based on this threshold, half of the acquired patents are associated with a Big Tech core technology.

[14] Out of the 5,056 inventors from firms acquired by Big Tech between 1996 and 2021, 90 are only listed in Reissue or Design patents, which are not associated with information on the patent's CPC classification. So the variable $\overline{Core_i}$ can be constructed for 4,966 inventors.

[15] See Appendix B.3 for statistics on *Core* aggregated at the target level.

Market value

We construct a proxy for market value based on three indicators of a patent's economic value: family size, grant lag and number of claims. The definitions of these variables and the reasoning behind their inclusion in the market value index are presented in Appendix B.4.

To construct the market value index, we propose two approaches. First, we define a simple linear combination of the three indicators of a patent's value. By taking the average value[16] of this linear combination for all the patents p filed by a given inventor i, we can define a first measure of the market value of the inventor's innovation history:

$$\overline{MarketVal_i} = \frac{1}{P}\sum_p \{FamilySize_p - GrantLag_p + AdjClaims_p\} \tag{a}$$

This index is further normalized such as to be comprised between 0 and 1 by min/max scaling.[17]

To improve on this first definition, we propose a second approach. We define a vector that captures as much as possible of the variation in the data along the three indicators of a patent's value. We thus intend to approximate our 3-D value space by a linear combination of all 3 (normalized) indicators along which the spread of patents is maximised. This vector is computed using a *Principal Components Analysis*, as described in Appendix B.5. By taking the total value of this vector for all the patents p filed by a given inventor i, we can define a second measure of the market value of the inventor's innovation history:

$$MarketVal_i = \sum_p \{ev^1 \widetilde{FamilySize}_p + ev^2 \widetilde{GrantLag}_p + ev^3 \widetilde{AdjClaims}_p\} \tag{b.1}$$

where ev^x are the coordinates of the vector.

[16] In this simple definition of the index, the three indicators of a patent's value have not been normalised, so indicators with higher expected values have more weight. The indicator with the highest weight is the family size: $E[FamilySize] > E[GrantLag] > E[AdjClaims]$. Aggregating by taking the mean value across all the patents filed by a given inventor ensures that we are not artificially boosting the index for inventors associated with proportionally higher levels of the family size variable.

[17] Min/max scaling: $\widetilde{X} = (X - min(X))\,/(max(X) - min(X))$.

In other words, the value of the inventor's innovation history is computed as the weighted sum of all the patents filed by this inventor pre-acquisition, where the weights capture these patents' market values. Because this definition allows to maximise the variation in our data, i.e. to best differentiate inventors based on their innovation history, we will use it as the baseline definition of the market value index.

Two important methodological choices have been made to construct this index. First, the summation over all the patents filed by a given inventor means that prolific inventors (i.e. filing many patents) will be on average associated with a higher market value. The reasoning behind this choice is that the more patents over which an inventor can be observed, the more information the acquirer can obtain about her innovation history. However, this choice implies a size effect that inflates the market value index for more productive inventors. For this reason, we also propose an alternative index by taking the *average* value of the vector for all the patents filed by a given inventor:

$$MarketVal_i^{bis} = \frac{1}{P}\sum_p \{ev^1 \widetilde{FamilySize}_p + ev^2 \widetilde{GrantLag}_p + ev^3 \widetilde{AdjClaims}_p\} \tag{b.2}$$

A second methodological choice in the definition of our market value index is the exclusion of a variable capturing the technological importance of a patent for the development of subsequent technologies; the number of citations it receives (forward citations). This is because forward citations are a stock that builds over time, so patents published at different times cannot be compared. To overcome this problem, we propose to consider the number of forward citations received by a patent *over a period of five years* after its publication date. This means that this variable can only be used for patents published at least five years before the end of our study period (in July 2022), so in July 2017. Because publication typically occurs around 18 months after the filing date (Squicciarini, Dernis, and Criscuolo 2013), we restrict this third specification of our market value index to patents filed - and hence inventors acquired - before January 2016. Forward citations are added to the market value index linearly:

$$MarketVal_i^{ter} = \sum_p \{ev^{t1}\widetilde{FamilySize_p} + ev^{t2}\widetilde{GrantLag_p} + ev^{t3}\widetilde{AdjClaims_p} + ev^{t4}\widetilde{FwdCit_p}\} \quad \text{(b.3)}$$

Talent acquisition

Our main variable of interest captures whether an inventor keeps patenting for her acquirer after acquisition. We collect information on the unique inventor id(s) in the target's patent portfolio, and check whether these same id(s) can be found in the acquirer's portfolio after acquisition: $Talent_i = 1$ if an inventor's id i in the target's patent portfolio can be found in the acquirer's portfolio.[18] This allows us to identify inventors who keep patenting for their acquirer. Those inventors who don't ($Talent_i = 0$) could have either stopped patenting altogether after the acquisition, or they could be patenting for some other firm; in both cases, they are not developing intellectual property for their acquirer.[19]

This measure can be further refined to better capture potential synergies deriving from acquisition by focusing on cases where newly hired inventors collaborate with the acquirer's staff: $Talent_i^* = 1$ if the acquirer's portfolio contains a patent associated with an inventor's id i from the target's portfolio together with an inventor's id from a patent published by the acquirer pre-acquisition.

From Appendix B.6, we can see that, in 74% of cases, at least one start-up inventor will be further patenting for their acquirer and, in 49% of cases, they do so in collaboration with some of their acquirer's existing employees. After normalizing by the total number of inventors, we can also see that inventors hired at Facebook's are proportionally more likely to keep developing patented technologies than inventors hired at Google's

[18] Some inventors (< 1% sample) filed patents for different targets pre-acquisition. For these inventors, we consider only their last employer (i.e. the target for which they last filed a patent) before acquisition.

[19] Let us note that using patent data to identify target inventors suffers from a limitation; it only captures inventors who publish patents. Some technologies might not be patented, because they are simply not patentable, or due to a low probability of imitation and/or high costs of patenting.

or Microsoft's.

2.2.3 Working sample

We describe in Table 2.2 below the profiles of target inventors at the time of their acquisition by a tech giant, separately for inventors who will keep innovating for their acquirer (i.e. who will file some patent for their acquirer after acquisition) and those who won't.

Table 2.2: Statistics over all the patents filed by Big Tech-acquired inventors

	Stop innovating 3293 (65.13%)	Keep innovating 1763 (34.87%)	Test
Patents (Total)	8559	6950	.
Patents over last 9m (Total)	714	1388	.
Patents (Mean)	2.60 (5.59)	3.94 (6.12)	<0.00
Patents over last 9m (Mean)	1.33 (0.74)	1.98 (2.86)	<0.00
First author patents (%)	0.55 (0.46)	0.58 (0.43)	0.05
First author patents over last 9m (%)	0.47 (0.48)	0.54 (0.46)	0.01
Core patents (%)	0.44 (0.48)	0.49 (0.49)	<0.00
Core patents over last 9m (%)	0.57 (0.49)	0.53 (0.49)	0.14
Market value (Mean)	-0.04 (0.08)	-0.05 (0.18)	0.18
Market value over last 9m (Mean)	-0.04 (0.08)	-0.05 (0.19)	0.02
US patents (%)	0.83 (0.38)	0.79 (0.41)	<0.00
US patents over last 9m (%)	0.82 (0.39)	0.74 (0.44)	0.00
Co-authored patents (%)	0.94 (0.22)	0.94 (0.20)	0.78
Co-authored patents over last 9m (%)	0.95 (0.21)	0.95 (0.21)	0.79
Months between first filing and acquisition	52.02 (38.06)	40.33 (32.10)	<0.00
Months between last filing and acquisition	40.00 (34.31)	22.21 (22.74)	<0.00

Mean (Standard deviation): p-value from a pooled t-test.
Frequency (Percent%): p-value from Pearson test.
This table presents patent statistics at the inventor-level. Negative average Market values come from the fact that patents filed by Big Tech-acquired inventors are less valuable than the average patent in the USPTO database. However, acquired patents are different in many respects from the average patent. The market value index should not be used to compare inventors across the whole patent database. Instead, we will be using it to compare acquired inventors among themselves.

Out of the 5,056 inventors from firms acquired by Big Tech between 1996 and 2021, 1,763 (35%) will file some patent for their acquirer after acquisition. Interestingly, we observe that inventors who keep innovating for their acquirer had filed on average more patents pre-acquisition, so they were on average more prolific than inventors who stop patenting after acquisition. They were also more often identified as "first author" of the patent, indicating a lead over the research project, and had filed, on average, more patents originating from outside the United States. In addition, these inventors who keep patenting for their acquirer appear to have started filing patents for the target later in time, so they could be considered more "junior" inventors. Finally, they tend to file their last patent for the target closer to the acquisition date, which is a sign that they were still active short before acquisition.

2.3 Innovating for the acquirer

In this section, we will examine the question of whether and why inventors of Big Tech core technologies (i.e. who have filed some "core" patents) are more likely to further innovate for their acquirer. To do so, we want to isolate the effect of working on a core technology from other potential determinants of whether a target inventor keeps innovating for her acquirer.

We consider the following inventor's characteristics (at the time of acquisition): the number of patents filed by the inventor for the target, whether the inventor has filed some patent as first author - since "lead" engineers are more likely to have been the targets of the acquirer - and whether the inventor has filed some patent originating from the US. We also include the inventor's first patent filing - as a proxy for how experienced she is - and her last filing before acquisition - which captures whether the inventor was still active at the time of acquisition. Finally, the outcomes of interest ($Talent$ and $Talent^*$) are constructed based on whether an inventor was *ever* found in her acquirer patent portfolio after acquisition. The period over which these outcomes are measured thus varies depending on the time at which the inventor joins the acquirer's team (i.e. the acquisition date). For this reason, we must control for the length of the period over which the outcome is measured; the number of months between acquisition and the end of the study period (in July 2022). In a robustness check, we will propose a more conservative approach with regard to the comparability of the outcome across inventors by defining a fixed study period.

First, we test whether the technological environment in which Big Tech acquisitions take place influences talent engagement (Section 2.3.1). We find that inventors working pre-acquisition on technologies in which their acquirer has been focusing its own innovative efforts are more likely to keep innovating for their acquirer after acquisition. Second, we will evaluate two hypotheses that could explain this finding: i. Big Tech could more easily assess the value of inventors when they work on its core technologies (Section 2.4.1), ii. Big Tech could be acquiring talent with the aim to have more people working on these core technologies in its research team (Section 2.4.2).

2.3.1 Baseline - Inventors of Big Tech core technologies

In this section, we want to explore the role of Big Tech technology leadership on the likelihood for acqui-hired inventors to further innovate for their acquirer. Studies such as Argentesi et al. (2021) and Gautier and Lamesch (2021) suggest that acquirers focus on reinforcing their existing segments. Conversely, Big Tech may acquire smaller firms to leverage specialized expertise and talent for expansion into peripheral activities (Varian 2021). With this baseline analysis, we aim to test empirically whether inventors of technology fields in which Big Tech is already strong are more likely to further innovate for their acquirer.

We consider all inventors of Big Tech-acquired technologies. Among these target inventors, we identify those who keep innovating for their acquirer from the patents they file post-acquisition. The outcome variable $Talent_i = 1$ when the inventor i's id can be found in her acquire's patent portfolio after acquisition, 0 otherwise. The probability for a start-up inventor to file some patent for her acquirer is modelled as follows:

$$P(Talent_i = 1) = F\left(\alpha + \beta Core_i + \tau MSinceAcqui_i + \sum_k \gamma_k X_{i,k} + \sum_l \zeta_l Y_{i,l} + \mu BT_i\right) \quad (2.1)$$

where $F(z) = \frac{e^z}{(1+e^z)}$ is the cumulative logistic distribution.

The regressor of interest ($Core_i \in \{0, 1\}$) captures whether inventor i has filed some patent belonging to one of her acquirer's core technology fields (i.e. fields of which the acquirer owns at least 1% of all published patents) by the time of acquisition. $MSinceAcqui_i$

captures the number of months between acquisition and the end of the study period. Inventor-level controls X_i include the number of patents filed by the inventor for the target ($PatentsCount_i$), the year in which the inventor filed a patent for the first time ($FirstFiling_i$), whether the inventor filed a patent originating from the United States ($US_i \in \{0,1\}$), whether the inventor filed some patent as first author ($FirstAuthor_i \in \{0,1\}$), and the number of months since the inventor's last patent filing for the target ($MSinceLastFil_i$). We also include the acquirer's identity fixed effects ($BT_i \in \{GOOG, APPL, FCBK, AMZN, MSFT\}$). Target-level controls Y_i include the firm's incorporation year ($Incorp_i^T$), country of origin ($US_i^T \in \{0,1\}$), number of employees ($FirmSize_i^T \in \{<50,\ [50;250[, \geq 250\}$), and funding amount category ($Funding_i^T \in \{\leq\ \$230M;\ ]\$230M, \$1,730M];\ > \$1,730M\}$).

Estimates

The parameters of Model (2.1) estimated by Maximum Likelihood can be found in Table 2.3 when the outcome variable is defined based on all target inventors found in the acquirer's portfolio after acquisition ($Talent_i$), and in Appendix B.7 when the outcome variable only considers those target inventors collaborating with some of their acquirer's employees after acquisition ($Talent_i^*$). We see that having filed a core patent is positively associated with the probability of further patenting for the acquirer.

The β parameters can be more easily interpreted in odds ratios (e^{β}). If we look at the most complete model in Table 2.3 (last column), we have $e^{\hat{\beta}} = 1.76^{***}$; the odds of further patenting for the acquirer are about 76% greater for inventors active in their acquirer's core technology fields at acquisition. From Appendix B.7, we see that these inventors are also on average more likely ($e^{\hat{\beta}} = 1.59^{*}$) to further patent for their acquirer in collaboration with some of the acquirer's existing employees.

Table 2.3: Inventors innovating for their acquirer, Model (2.1)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core	0.404***	0.356***	0.313***	0.291***	0.292***	0.452***	0.445***	0.510***	0.563***
	(4.82)	(4.05)	(3.50)	(3.19)	(3.19)	(3.82)	(3.73)	(3.54)	(3.66)
MSinceAcqui	0.012***	0.029***	0.029***	0.012***	0.013***	0.017***	0.017***	0.019***	0.020***
	(15.92)	(19.93)	(19.02)	(6.72)	(7.11)	(7.83)	(7.83)	(7.97)	(8.04)
PatentsCount	0.070***	0.107***	0.105***	0.040***	0.034***	0.035***	0.036***	0.032***	0.032***
	(9.12)	(11.88)	(11.57)	(4.66)	(4.09)	(3.73)	(3.75)	(3.41)	(3.37)
FirstFiling		0.189***	0.187***	-0.022	-0.011	-0.018	-0.018	-0.034	-0.038
		(14.52)	(13.55)	(-1.14)	(-0.56)	(-0.83)	(-0.83)	(-1.40)	(-1.56)
US			-0.260***	-0.227**	-0.227**	-0.264**	-0.014	-0.106	-0.610
			(-3.02)	(-2.57)	(-2.56)	(-2.08)	(-0.03)	(-0.24)	(-1.09)
MSinceLastFil				-0.030***	-0.029***	-0.030***	-0.030***	-0.034***	-0.036***
				(-14.17)	(-13.74)	(-12.79)	(-12.78)	(-12.93)	(-13.08)
FirstAuthor					0.294***	0.334***	0.334***	0.272***	0.266***
					(4.21)	(4.31)	(4.32)	(3.23)	(3.10)
$Incorp^T$						0.009	0.009	-0.014	-0.018*
						(1.04)	(1.02)	(-1.34)	(-1.65)
US^T							-0.255	-0.064	0.358
							(-0.67)	(-0.15)	(0.68)
$FirmSize^T$								-0.568***	-0.611***
								(-7.31)	(-5.28)
$Funding^T$									0.082
									(0.67)
Constant	-1.779***	-383.650***	-378.974***	43.302	20.953	15.684	15.850	94.411*	112.447**
	(-14.66)	(-14.59)	(-13.61)	(1.11)	(0.53)	(0.35)	(0.35)	(1.89)	(2.19)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

2.3.2 Robustness checks

To test for the robustness of our results, we replicate our baseline analysis with alternative definitions of the dependent (*Talent*) and independent (*Core*) variables. The parameter of Model (2.1) estimated based on this alternative specification are presented in Appendix B.8.

Alternative definitions of the dependent variable

Fixed observation period To better ensure the comparability of the outcome of interest (*Talent* or *Talent**) across all inventors, we propose to construct the analysis based on a fixed time period. The idea is that the period over which we test whether a target inventor is found in her acquirer patent portfolio would be the same for all inventors.

We choose to define $Talent_i = 1$ when the inventor i's id can be found in her acquire's patent portfolio within 4 years after acquisition ($Talent_i = 0$ if she never patents for her acquirer *or* if she does so more than 4 years after acquisition). We thus ignore all inventors who were acquired less than 4 years before the end of the data collection in July 2022. So we consider inventors acquired up until July 2018.[20]

Just like for the baseline estimates, we find that inventors active in their acquirer's core technology fields are on average more likely to further patent for their acquirer, and the effect is even larger: the odds of further patenting for the acquirer are about 2.61 times greater for inventors active in their acquirer's core technology fields (see Table B.6).

Two months buffer after acquisition As a second robustness check, we want to accommodate a period during which uncertainty persists regarding the patented technology investigator (whether still target or already acquirer).

We choose to define $Talent_i = 1$ when the inventor i's id can be found in her acquire's patent portfolio from two months onwards after acquisition ($Talent_i = 0$ if she never patents for her acquirer *or* if she does so less than two months after acquisition).

On this basis, we observe that inventors active in their acquirer's core technology fields are 2.16 time more likely to further patent for their acquirer (see Table B.8). So our previous results are also robust to this alternative specification.

Inventors filing under the acquirer's name So far, we have included in the acquirer's portfolio all the patents filed under the acquirer's name *or* under the target's name after

[20] This 4 years threshold also allows to account for an observed drop in the probability to further patent for the acquirer when acquired after 2018 (see Appendix B.9). This can partially be explained by the fact that the USPTO PatentsView data covers granted patents, so the most recent patents are less likely to be observed simply because they might not be granted yet.

acquisition. This means that an inventor who keeps filing patents under her target's name after acquisition is considered as patenting for her acquirer. Of all the inventors associated with $Talent_i = 1$ in our baseline analysis, 19% have not filed any patent under their acquire's name.

One could argue that inventors who don't use their acquirer's name when filing a patent after acquisition are not fully integrated whithin their acquirer's research lab. Instead, they stay somehow more independent and this might go along with a research history that is more distant from their acquirer's core fields, which might drive our β estimates downwards. For this reason, we propose to check for the robustness of our results to restricting $Talent_i = 1$ to those target inventors filing some patents under their acquirer's name.

The outcome of this alternative specification confirms our intuition, as the β estimates appear even larger than in the baseline analysis: inventors active in their acquirer's core technology fields are on average 2.93 times more likely to further patent under their acquirer's name after acquisition (see Table B.10).

Alternative definitions of the independent variable

Inventors of core technologies within a year As an additional robustness check, we want to test whether our results hold when using an alternative definition of the regressor of interest; instead of capturing whether an inventor has *ever* filed some patent belonging to one of her acquirer's core technologies, $Core_i \in \{0, 1\}$ will now capture whether she does so *within a year* before acquisition.

Our results show that the coefficient associated with $Core_i$ is still positive, but it is no longer statistically significant for the more complete model specifications (see last columns of the regression output in Table B.12). This absence of statistical significance could be explained by the reduced sample size; only 852 inventors (17% of inventors from Big Tech targets) have been active in their acquirer's core technology fields within a year of their acquisition. So statistical power might not be sufficient to estimate the most complete versions of our model (i.e. including all control variables).

Core technologies defined at the 3-digit level We define Big Tech's core technologies based on the relative frequency at which those technologies appear in Big Tech's patent portfolios. For our baseline analysis, we computed the number of patents classified in each cpc subsection at the two-digits level (CPC section symbol followed by a two-digit number, e.g. G06 for 'Computing; Calculating or Counting'). With this robustness check, we propose a more precise mapping by narrowing the definition to the three-digit level (e.g. G06F for 'Electric Digital Data Processing'). Using this alternative definition of core technologies, the variable *Core* can be constructed in the same way as before (see Section 2.2.2).

In Table B.14, we see that our main results are robust to this alternative definition; inventors active in their acquirer's core (3-digit) technology fields are 1.62 time more likely to further patent for their acquirer.

2.4 Extension: Assessing potential explanations for our baseline results

In the previous section, we found that inventors working pre-acquisition on technologies in which their acquirer is focusing its own innovative efforts are more likely to keep innovating for their acquirer after acquisition. In this section, we examine two potential explanations for this finding.

2.4.1 1^{st} hyp.: Inventors of core technologies are more easily assessed

Could Big Tech more easily assess the value of inventors when they work on its core technologies, and hence better exploit their talent, which would explain that these inventors are more likely to stay active after acquisition?

To examine this first hypothesis, we use the $MarketVal_i$ index to proxy for the value of inventor *i*'s innovation history (i.e. the market value of the patents this inventor filed pre-acquisition, as defined in Section 2.2.2). Next, we rewrite Model (2.1) to allow the effect of this value index on the probability for the inventor to keep innovating for her acquirer to vary depending on whether the inventor has filed some core patent ($Core_i$):

$$P(Talent_i = 1) = F(\alpha' + \beta_1 Core_i + \beta_2 MarketVal_i + \beta_3 Core_i MarketVal_i + \tau' MSinceAcqui_i + \sum_k \gamma_k' X_{i,k} + \sum_l \zeta_l' Y_{i,l} + \mu' BT_i) \quad (2.2)$$

For inventors associated with $Core_i = 1$, the odds ratio for a one-unit increase in the value of their innovation history is $e^{\beta_2+\beta_3}$, for those with $Core_i = 0$ it is e^{β_2}.

Estimates

The parameter estimates of Model (2.2) are presented in Table 2.4. Based on the baseline definition of the $MarketVal_i$ index (see equation (b.1)), we estimate that, for inventors active in technologies outside of their acquirer's core business ($Core_i = 0$), a history of more valuable patents does not translate into a higher probability to keep innovating for their acquirer; they actually have 97% fewer chances to do so ($e^{\hat{\beta}_2} = .03^{***}$) for each unit increase in their innovation value.

The coefficient associated with $MarketVal_i$ captures two mechanisms: inventors who stay at their acquirer's are expected to be more likely to further patent if they have a more valuable innovation history, but a more valuable innovation history might also help them find another job more easily at some other firm's. Based on our results, the latter mechanism seems to be stronger for non-core inventors. In Appendix B.10, we can see that these results are robust to using the alternative definitions of the market value index (see equations (a), (b.2) and (b.3)).

These findings suggest that talented inventors active in technologies outside of their acquirer's core business might not be recognized for their true worth at their acquirer's. Instead, a history of more valuable patents might help them find another job in which their talent could be better exploited. This corroborates the hypothesis according to which being better assessed by their acquirer partly explains why inventors of core technologies are more likely to keep innovating for their acquirer.

Table 2.4: Inventors innovating for their acquirer, Model (2.2)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core=1	0.487***	0.465***	0.433***	0.381***	0.376***	0.564***	0.559***	0.612***	0.687***
	(5.57)	(5.05)	(4.63)	(4.01)	(3.95)	(4.55)	(4.49)	(4.11)	(4.30)
MarketVal	-1.866***	-3.427***	-3.632***	-2.592***	-2.424***	-2.561***	-2.560***	-3.223***	-3.358***
	(-3.83)	(-6.47)	(-6.76)	(-4.84)	(-4.52)	(-4.33)	(-4.33)	(-5.11)	(-5.23)
Core=1 × MarketVal	2.338***	3.157***	3.436***	2.843***	2.717***	3.395***	3.392***	3.662***	3.860***
	(4.10)	(5.21)	(5.57)	(4.69)	(4.48)	(4.97)	(4.97)	(5.18)	(5.37)
MSinceAcqui	0.011***	0.030***	0.031***	0.014***	0.014***	0.018***	0.018***	0.021***	0.022***
	(15.40)	(20.24)	(19.46)	(7.30)	(7.58)	(8.09)	(8.09)	(8.46)	(8.54)
PatentsCount	0.065***	0.101***	0.099***	0.037***	0.032***	0.031***	0.032***	0.030***	0.030***
	(8.40)	(11.25)	(10.90)	(4.34)	(3.83)	(3.38)	(3.39)	(3.22)	(3.21)
FirstFiling		0.210***	0.212***	0.002	0.010	0.001	0.001	0.004	0.004
		(15.29)	(14.49)	(0.09)	(0.48)	(0.04)	(0.04)	(0.15)	(0.14)
US			-0.257***	-0.218**	-0.217**	-0.262**	-0.091	-0.245	-0.731
			(-2.94)	(-2.43)	(-2.42)	(-2.06)	(-0.23)	(-0.56)	(-1.30)
MSinceLastFil				-0.029***	-0.028***	-0.029***	-0.029***	-0.033***	-0.034***
				(-13.54)	(-13.20)	(-12.45)	(-12.44)	(-12.24)	(-12.34)
FirstAuthor					0.260***	0.299***	0.300***	0.224***	0.214**
					(3.67)	(3.82)	(3.82)	(2.62)	(2.45)
Incorp^T						0.012	0.012	-0.012	-0.018
						(1.27)	(1.26)	(-1.13)	(-1.60)
US^T							-0.174	0.029	0.485
							(-0.46)	(0.07)	(0.91)
FirmSize^T								-0.589***	-0.597***
								(-7.48)	(-5.08)
Funding^T									0.039
									(0.32)
Constant	-1.744***	-427.012***	-429.059***	-4.857	-21.036	-27.318	-27.166	15.161	27.368
	(-13.84)	(-15.35)	(-14.54)	(-0.12)	(-0.51)	(-0.57)	(-0.57)	(0.29)	(0.50)
Observations	4880	4880	4720	4720	4720	3967	3967	3509	3435
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A question that naturally arises from this result is whether these inventors of valuable core technologies are expected to keep working in similar technology fields after acquisition. Is Big Tech hiring experts because it wants to develop their areas of expertise, or is it hiring the best talents, regardless of their specialization, to further expand into peripheral fields? This question will be examined in the next section.

2.4.2 2^{nd} hyp.: Big Tech acquires talent to improve its core technologies

We have argued so far that inventors of core technologies can be better assessed by their potential acquirer, which could partially explain that they are more likely to keep innovating after acquisition. But uncertainty persists as to what Big Tech ends up using this talent for once it has joined the new research team. Are they simply more often further innovating for their acquirer because they were better assessed in the first place, or could it be a strategy for Big Tech to acquire inventors skilled in its core technologies with the aim to have more people working on these technologies in its research team? In other words, are these inventors of core technologies hired for their general talent, or for their expertise in specific fields?

To examine this question, we want to assess whether the best inventors of Big Tech core technologies keep working on these technologies after they joined their acquirer's team. We thus focus on inventors who are further innovating for their acquirer (i.e. who have filed some patent for the target before being acquired *and* for their acquirer after acquisition, $Talent_i = 1$). This means that our analysis is prone to a selection bias: the technology choices of inventors who do not keep patenting for their acquirer are not observed, while patenting for the acquirer might not be random. For instance, those inventors with a more positive relation between their innovation value and the probability to file core patents could be less likely to keep patenting for their acquirer and hence to be observed.

Inventor keeps innovating for the acquirer

Inventor focuses on core technologies   0

1   0

So there might be two dimensions to the relationship between an inventor's innovation value and her propensity to keep working on her acquirer's core technologies. First, the direct effect of the inventor's ability to develop certain technologies. Second, the fact that some inventors might be *selectively* less likely to be observed after acquisition (e.g. if they don't comply with what is expected of them based on their value type). To study the determinants of how inventors direct their innovative effort while keeping constant the effect of selection into further innovating for the acquirer, we propose to use a Heckman (1979)-type selection model.

We collect information on how close the patents filed by target inventors are to their acquirer's technology. In Section 2.2.2, we defined the $\overline{Core_i}$ variable as the share of fields listed in patents filed by inventor i that belong to her acquirer's core business. This variable can be constructed with a time dimension based on the patents filing dates; up to two years before the acquisition of the target, and over a period of two years starting two months after acquisition. This two months buffer aims to accommodate a period during which uncertainty persists regarding the patented technology investigator. We obtain a dataset at the inventor-level with information on the share of core fields before acquisition ($\overline{Core_i^{Before}}$), for both inventors who keep patenting for their acquirer and those who don't, and on the share of core fields over a period starting two months after acquisition ($\overline{Core_i^{After}}$), for inventors who keep patenting for their acquirer only.[21] [22]

Now, we want to assess whether target inventors who have patented more valuable inventions pre-acquisition are more likely to keep working on their acquirer's core technologies. The regression equation thus aims to capture the effect of the market value of the inventors' innovation history ($MarketVal_i$) on the share of core fields averaged across

[21]Out of the 2,645 inventors observed (up to 2 years) before acquisition, 1,037 inventors keep patenting for their acquirer (of which 1,013 are associated with some information on $MarketVal$ and all inventor-level control variables), coming from 134 targets, and 1,608 do not keep patenting for their acquirer (of which 1,501 are associated with some information on $MarketVal$ and all inventor-level control variables), coming from 144 targets.

[22]For inventors who keep patenting for their acquirer, we find no significant difference between the share of core fields before and after acquisition: based on a t-test, we cannot reject that the mean value of $\overline{Core_i}$ across all inventors filing before and after acquisition (.409 and .407, respectively) is the same. So, on average, we do not observe a shift in the probability for a target inventor to patent technologies belonging to her acquirer's core field.

all the patents filed by this inventor after acquisition ($\overline{Core_i}^{After}$):

$$\overline{Core_i}^{After} =$$
$$\alpha + \eta_1 MarketVal_i + \eta_2 Core_i^{Before} + \eta_3\, MarketVal_i\; Core_i^{Before} + \sum_k \gamma_k x_{i,k} + \mu BT_i + u_i \tag{2.3}$$

where x_i contains the year in which the inventor filed a patent for the first time ($FirstFiling_i$), whether the inventor filed a patent originating from the United States (US_i), the share of patents the inventor filed as first author ($FirstAuthorSh_i$), the number of months between acquisition and the end of the study period ($MSinceAcqui_i$), and $u_i \sim N(0, \sigma)$. $Core_i^{Before}$ accounts for the fact that high levels of $\overline{Core_i}^{After}$ could be associated with two different scenarios: an inventor starts patenting in her acquirer's core fields after acquisition ($Core_i^{Before} = 0$), or she keeps patenting in these same fields she was already working on before acquisition ($Core_i^{Before} = 1$). Since we want to explain why core inventors are more likely to further innovate for their acquirer after acquisition, we are interested in the case where $Core_i^{Before} = 1$.

Selection model

In order to correct for selection, we need to define an exclusion restriction. We choose the variable associated with the number of months since the inventor's last patent filing for the target ($MSinceLastFil_i$), which is a proxy for whether the inventor is still active at the time of acquisition.

Exogeneity The instrument we propose is expected to be uncorrelated with u_i: being active at the time of acquisition should not be directly linked with the share of core fields in the patents an inventor files after acquisition. It is difficult to test for the validity of this exclusion restriction based on $\overline{Core_i}^{After}$, because the effect of $MSinceLastFil_i$ could be direct, but also indirect through the probability of being observed after acquisition ($Talent_i = 1$). Hence, we propose to test it based on $\overline{Core_i}^{Before}$.

We estimate the Pearson (1896) product-moment correlation coefficient between

$\overline{Core}_i^{Before}$ and $MSinceLastFil_i$ on the full sample of inventors who keep patenting for their acquirer and those who don't. We obtain a correlation coefficient of -0.008 that is not different from zero at all standard confidence levels. We conclude that the correlation between our instrument and the share of core fields in the patents an inventor files *before* acquisition is not significant. This suggest that the direct effect of this instrument on the share of core fields *after* acquisition should also be insignificant, and that only an indirect effect through the probability of being observed after acquisition would remain.

Relevance The instrument we propose could also be correlated with the inventor-level controls (in x_i and BT_i), so these controls may influence $Talent_i$ via $MSinceLastFil_i$. To isolate the direct effect of $MSinceLastFil_i$ on $Talent_i$ from the effect these controls may have on $Talent_i$ through $MSinceLastFil_i$, we include x_i and BT_i in the selection equation, of which we define a Probit estimate:

$$P(Talent_i = 1) = \Phi\,(\alpha^1 + \eta_1^1 MarketVal_i \ + \eta_2^1 Core_i^{Before} + \eta_3^1\ MarketVal_i\ Core_i^{Before} \\ + \sum_k \gamma_k^1 x_{i,k} + \mu^1 BT_i + \kappa\ MSinceLastFil_i) \quad (2.4)$$

where Φ is the standard cumulative normal distribution.

To be relevant, the instrument must be correlated with the probability to be observed after acquisition; the κ parameter from this selection equation must be significantly different from zero. We can check this from Appendix B.11, presenting the first-step parameter estimates from equation (2.4). We confirm the relevance of our instrument based on the significance of the $\hat{\kappa}$ parameter ($-.01^{***}$).

Estimates

Using Heckman's (1979) procedure, we estimate the two-step parameters from the regression equation (2.3) augmented with the nonselection hazard computed from the fitted values of the selection equation (2.4):

$$\overline{Core}_i^{After} = \alpha^2 + \eta_1^2 MarketVal_i + \eta_2^2 Core_i^{Before} + \eta_3^2\, MarketVal_i\, Core_i^{Before} + \sum_k \gamma_k^2 x_{i,k} + \mu^2 BT_i + \lambda \frac{\phi(P(\widehat{Talent_i = 1}))}{\Phi(P(\widehat{Talent_i = 1}))} + v_i \quad (2.5)$$

where ϕ and Φ are the standard normal probability density and cumulative distribution functions, respectively.

Provided that $MSinceLastFil$ is a valid exclusion restriction, which we argued in the previous section, these parameters capture the responses in the observed values of $\overline{Core}^{After}$ to independent variations in the regressors, i.e. keeping constant the effect of selection into $Talent = 1$.

Table 2.5 presents the two-step parameter estimates from model (2.5), where the parameter associated with $MarketVal$ thus captures the share of the variation in $\overline{Core}^{After}$ explained by the exogenous part of the market value of the inventor's innovation history. The two-step parameter estimates of $\eta_1^2 + \eta_3^2$ is negative $(-.18^{***})$, so core inventors who have patented more valuable inventions pre-acquisition are *less* likely to keep innovating in their acquirer's core technologies (and this is independent from whether they keep innovating for their acquirer at all).[23]

Next to our baseline definition of the independent variable, $MarketVal$, this model is also estimated based on the three alternative market value indices defined in equations (a), (b.2) and (b.3). As can be seen from Appendix B.12, our results are robust to these alternative definitions: $\hat{\eta}_1^2 + \hat{\eta}_3^2 = -.29^{**}$ (for $\overline{MarketVal}$), -1.76^{***} (for $MarketVal^{bis}$) and $-.16^{***}$ (for $MarketVal^{ter}$).

[23] We also observe that the λ parameter associated with the nonselection hazard in the second step of the selection model is not significant. This means that there does not seem to be a bias due to, for instance, inventors associated with both a high innovation history market value and many core patents not being observed.

Table 2.5: Heckman two-step parameters

MeanCore_After	
MarketVal	0.034
	(0.093)
Core_Before = 1	0.654***
	(0.022)
Core_Before=1 × MarketVal	-0.216*
	(0.120)
MSinceAcqui	-0.000
	(0.000)
FirstFiling	0.004
	(0.003)
US	-0.080***
	(0.018)
FirstAuthorSh	0.042**
	(0.018)
Constant	-8.926
	(5.757)
lambda	-0.022
	(0.068)
Observations	2,515
BT dummies	Yes

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Our results thus show the opposite of what would be expected if Big Tech were acquiring inventors skilled in its core technologies to increase the number of in-house researchers focused on these areas. Instead, core inventors with a history of more valuable patents appear to direct their innovation efforts towards novel research projects.

2.5 Conclusion

In this paper, we analyse the innovative activity of inventors from firms acquired by Big Tech. Drawing on firm-level and patents databases, we find that inventors working pre-acquisition on technologies in which their acquirer is focusing its own innovative efforts are more likely to keep filing patents for their acquirer after acquisition.

As we show, this outcome can be partially attributed to Big Tech's enhanced talent assessment capabilities in its core technology fields, which explains that inventors of core technologies are more likely to stay active after acquisition. However, we also find that, once acquired, these talented core inventors are redirected towards fields in which their acquirer has less expertise. These results suggest that Big Tech acquires inventors in its core technology fields mainly because it can better assess their value, rather than to maintain their focus on those core technologies.

This paper highlights the interplay between talent acquisition and technological specialisation in the digital economy. Our analysis focuses on inventors joining their acquirer's team. An interesting extension for future research would be to investigate what happens to inventors who leave their company after its acquisition by a tech giant and join some other firms.

Chapter 3

How and Where Does Big Tech Disrupt?

Abstract

Despite their considerable R&D investments, the tech giants might not always seem to have what it takes to bring disruptive innovation to the market. For these leading firms, the acquisition of start-ups could be a way to disrupt at lower cost. But acquiring and shelving disruptive inventions might also enable big tech platforms to stifle disruption in markets where they are already strong. We propose to use a patent-based measure of the extent to which a firm develops disruptive innovation. We first document, based on this measure of disruption, Big Tech's innovative activity. Alongside the 'buy vs build' principle, we find that Big Tech's acquired inventions tend to be more disruptive than those that it develops internally. Next, we test whether the adoption of those acquired disruptive technologies varies with the competitive environment in which the acquisition takes place. Our results suggest that disruptive technologies that have been adopted by Big Tech belong to industries in which these tech giants hold a weaker market position.

3.1 Introduction

When studying innovation, the economic literature has traditionally mainly focused on *how much* innovation is produced. More recently, an increasing number of studies have been focusing on the *type* of innovation that is produced. Is it some type of innovation that just marginally improves an existing technology, or does it completely revolutionise the state of the art?

Since Christensen (1997), the literature refers to these innovation types as "sustaining" versus "disruptive" innovation. Gans (2016) developed a more modern notion of disruption, from the perspective of the incumbent: *"when successful firms fail because they continue to make the choices that drove their success"* (p.9). Another influential definition, proposed by Cabral (2023), distinguishes between an innovation that establishes a new dominant firm (competition for the market), as opposed to an innovation that results in a higher technology level (competition in the market). In line with this idea, Acemoglu, Akcigit, and Celik (2014) refer to technologies causing the most fundamental "creative destruction". In the management literature, Govindarajan and Kopalle (2006) developed a scale to measure disruption based on five characteristics of the innovation,[1] with a focus on how the market only gradually comes to value the innovation's attributes.

In this paper, we will prefer a less market-oriented and more technology-oriented definition of disruption. To define a technological change, the literature has long been using the distinction between *improving* an existing technology and *introducing a new approach* to technical practice (Reinganum 1983; Tushman and P. Anderson 2018). According to this definition, disruption identifies, from the mass of minor improvements on existing technologies, those few novel technologies that significantly impact technical progress (Scherer and Harhoff 2000; Trajtenberg 1990).

We apply our analysis of disruptive technologies to digital markets. To do so, we focus on the leading firms of the sector in the US, often grouped under the labels Big

[1]The five characteristics used by Govindarajan and Kopalle (2006) to define a disruptive innovation are: i. it underperforms on dimensions mainstream customers value, ii. mainstream customers initially do not value the innovation's attributes, iii. it is simpler and cheaper than existing offerings, iv. it appeals to low-end, financially unattractive niche customers, v. it can disrupt high-end markets with radical technology. Govindarajan and Kopalle (2006) tested for the internal validity of this scale based on a survey of senior executives from different industry sectors, in which they were asked to rank the innovations that had been commercialized by their strategic business unit (SBU) along the five items. For instance, to capture the ability of an innovation to attract new customer segments, the authors used the following question (to be answered on a 7-point scale, strongly disagree/strongly agree): "During the past 5 years, the new products that were introduced by this SBU were very attractive to a different customer segment at the time of product introduction.". They showed that executives belonging to the same SBU tend to give similar rankings.

Tech, GAFAM or tech giants. Each of these leading firms established its dominance by commercialising a disruptive technology: search engines for Alphabet, personal computers for Apple, social networks for Meta, online shopping for Amazon, and operating systems for Microsoft (Lemley and Wansley 2024). Since then, big tech platforms have themselves not been displaced by the disruptive technology of another innovative firm. There has, however, been many cases of digital start-ups bringing some disruptive innovation to the market.

As an example, let us consider the case of Agawi, an app-streaming start-up whose technology enabled users to access apps on their smartphones without having to download them first (Efrati 2015). While this technology was originally used for in-app advertisements, it provided an alternative method for discovering content within apps, directing users away from Google's search engine and web browsers (Boyacıoğlu, Özdemir, and Karim 2024). In 2014, Google acquired Agawi (Lunden 2015).

Such case study is very informative about the dynamics of disruptive technologies in digital markets, and it offers a strong rationale for exploring disruption patterns on a more aggregated level. To do so, we use in this paper information found in the patents published or acquired by Big Tech. We follow Arts, Hou, and Gomez (2021) to construct a patent-based measure of disruption. These authors build a database of the keywords extracted from the title, abstract, and claims of all the patents filed at the USPTO between 1980 and 2018, and they propose to identify patents covering fundamentally new technologies based on the 'new' keywords (i.e. keywords appearing for the first time in the USPTO database) that they contain.

Based on this measure of disruption, we first propose some empirical evidence on the extent to which Big Tech adopts disruptive technologies, both through in-house innovation and through the acquisition of start-ups. Theory predicts that start-ups are more able/have more incentives to develop disruptive inventions (Acemoglu, Akcigit, and Celik 2014; Christensen 1997; R. M. Henderson and Clark 1990), and this is backed-up by empirical evidence showing that incumbents often lag behind in introducing disruptive product innovation or delay its diffusion (Antonio and Kanbach 2023; Ben-Slimane, Diridollou, and Hamadache 2020; Hynes and Elwell 2016). In line with these predictions, we find that big tech platforms tend to disrupt less than their targets.

Google's technology might have been less disruptive than Agawi's, but this does not necessarily stem from Google's failure to disrupt. Instead, this might be a strategic choice on the part of the tech giant to 'buy' instead of 'build' (Caffarra, Crawford, and Valletti 2020). In other words, the disruption effort of Big Tech and of its targets might be strategic substitutes and, instead of trying to replicate Agawi's technology, Google might choose to buy it. A question broadly discussed in such scenario (see for instance Fumagalli, Motta, and Tarantino 2020) is whether Google will adopt this disruptive technology or, instead, shelve it to avoid directing users away from Google's mobile web browsers.

While a case-by-case analysis would be necessary to answer this question for each Big Tech acquisition separately, in this paper we propose a methodology that allows to study broader patterns. To do so, information on the new keywords contained in a patent text can also be used. First, we identify whether a target had a disruptive technology before being acquired from the combination of new keywords in its patent portfolio. In turn, if its acquirer is adopting the technology, these new keywords combinations should appear in the acquirer's own patent portfolio. We find that, for about 3/4 of Big Tech acquisitions, the disruptive technology is later adopted by its acquirer.

Finaly, we expect that variations might arise across market environments in which the innovative activities take place. Theoretical literature suggests that large incumbents may not only have more to loose from adopting a disruptive invention in markets that they are leading (Arrow 1972, Holmes, Levine, and Schmitz Jr 2012), but they might also seek to hinder the disruption brought by innovative start-ups in those markets (Shapiro 2011). We aim to provide empirical evidence to complement these theoretical predictions. To do so, we use two different metrics of the market environment. First, we define a proxy for market size based on the number of companies with similar business characteristics to Big Tech targets. Second, we build an indicator of dominance based on market shares data. Our results suggest that the adoption of disruptive technologies from acquired start-ups does not vary across market sizes. However, we find that disruptive technologies that have been adopted by Big Tech (historically) belong to fields in which the tech giants have a weaker market position (today).

Background and Research Questions

There is a suspicion that the tech giants are not very good at disrupting (Ezrachi and Stucke 2022, R. M. Henderson and Clark 1990, Schmidt 2023). Because of their size, they might be little responsive to the evolution of the market and they might lack creativity (Crémer, Montjoye, and Schweitzer 2019). They might also find it organisationally challenging to invest in disruptive technologies (Christensen 1997; Bresnahan, Greenstein, and R. M. Henderson 2011).

To address these challenges and stay ahead in the technological race, these industry leaders often resort to acquiring smaller, more flexible start-ups (Bryan and Hovenkamp 2020a). These high growth young firms are identified as the main drivers of innovation in the most recent endogenous growth models (Acemoglu, Akcigit, Alp, et al. 2018, Akcigit and Kerr 2018). This is particularly important in digital markets, where start-ups play a major role in bringing disruptive innovations (Jorgenson 2001, Zucker, Darby, and Brewer 1998). Many of today's most popular and successful products started out with smaller companies that were later acquired by larger corporations (Cabral 2018), e.g. Microsoft/PowerPoint (1987), Google/YouTube (2006), and Facebook/Instagram (2012). The acquisition of innovative start-ups can thus come as a substitute to the incumbent's own in-house disruption effort in a "build vs. buy" scenario (Caffarra, Crawford, and Valletti 2020).[2]

[2] We thus focus on the incumbent's disruption effort. Although it is a very interesting topic, this paper does not explore the relationship between the startup's own disruption efforts and its acquisition potential. The literature on this topic tends to predict a negative effect of acquisition. Anticipating its acquisition, the start-up strategically distorts the direction of its innovation in order to maximise the acquisition rents (Dijk, José L Moraga-González, and Motchenkova 2021, Katz 2021). This leads to less radical innovation and lower quality (Cabral 2018, Katz 2021). In addition, the ability to innovate for buyout makes it less desirable for the start-up to develop a disruptive innovation that would allow it to replace the dominant firm. Cabral (2018) refers to this phenomenon as the "complacency effect".

Building on this literature, our first research question boils down to:

First research question *Are Big Tech's acquired inventions more disruptive than their in-house innovation?*

Next, we will be interested in the relationship between disruption and the market environment. Aghion et al. (2005)'s textbook finding is "an inverted-U relationship between product market competition and innovation.".[3] Again, we aim to redirect the focus from *how much* innovation is produced to *how disruptive* is innovation across different market environments.

For most authors, disruption is not only defined by the nature of the technology change (i.e. improving an existing technology vs. developing a substitute to this existing technology), but also by its ability to threaten the incumbent's competitive position in one of its key business areas (Adner 2002; Boyacıoğlu, Özdemir, and Karim 2024): "whether the disruptive technology is improving from below along a trajectory that will ultimately intersect with what the *market* [i.e. existing users] needs." (Christensen 1997, p.54). As such, the market environment is an important determinant of Big Tech's incentives to disrupt, both internally or through acquisition.

On the one hand, large incumbents are expected to be more disruptive when they target adjacent markets, e.g. Microsoft Bing challenging Google in search (Federico, Morton, and Shapiro 2020), as they might lack the incentives to disrupt markets that they are leading. On the other hand, disruption is less desirable for these leading firms when they have fewer firms to compete with, because they have more to loose from the "forgone rent" (Arrow 1972, Holmes, Levine, and Schmitz Jr 2012) and less to gain from business-stealing effects (Bryan and Hovenkamp 2020a; Federico, Morton, and Shapiro 2020). Instead, the acquisition of disruptive start-ups can be viewed as a tactic to stifle disruption in markets where the acquirer is already strong (Lemley and Wansley 2024; Shapiro 2011), leading to both the loss of a disruptive entrant and reduced competitive pressure on the incumbent, e.g. Facebook's acquisition of Instagram (Federico, Morton,

[3] See also Sutton (2001).

and Shapiro 2020).[4]

In line with these considerations, we formalise a second research question:

Second research question *How does Big Tech's disruptive technology adoption vary across market environments?*

3.2 Data

For our different analyses, we will be using information found in US patents filed by Big Tech itself and by Big Tech-acquired firms.

We identify firms acquired by Big Tech and the dates at which their acquisitions were announced from four different databases: Standard & Poor's CapIQ (2022), Geoff, Marshall and Parker (2021), Gautier and Lamesch (2021), and the US Patent and Trademarks Office (USPTO) Patent Assignment Dataset (2022). Patent data is collected from the USPTO Patent Views database (e.g. patent number, patent date, application identifier, publication author, CPC technology field and inventors' names[5]), and is matched to Arts, Hou, and Gomez (2021)'s patent text measures. We identify 859 firms acquired by Big Tech between January 1996 and January 2021, of which 252 have filed at least one patent application before being acquired.

3.2.1 Disruption

Patent data is often used in the literature to empirically capture the distinction between an incremental vs. disruptive innovation.

Some authors choose to use information on patent citations (Si and Chen 2020) and on 'development paths', i.e. chains of patents citing one another. For instance, the dis-

[4] On the start-ups' side, Cabral (2018) models that, since an increase in firm dominance implies a bigger profit from becoming the next dominant firm, the entrant's incentives to develop a disruptive innovation that allows it to become the next dominant firm are higher in a more concentrated market.

[5] Harmonized inventors' names can be found at https://patentsview.org/download/data-download-dictionary.

ruption index created by Funk and Owen-Smith (2017) measures the number of patents citing a focal patent, excluding those that share common citations with this focal patent. The idea is that a disruptive patent will be cited by subsequent work that is less likely to also reference its predecessors (M. Park, Leahey, and Funk 2023). In a study by Momeni and Rost (2016) on the photovoltaic industry, a technology is identified as disruptive if the associated patents are highly cited and belong to multiple development paths. In Cheng et al. (2017), interdisciplinary patent citations are used as a proxy for technological disruption, defined as the impact of a technology on existing technology development trends. However, the use of patent citations to measure disruption suffers from an important limitation; it does not reflect the technical content of the patent itself (Arts, Hou, and Gomez 2021).

Unlike patent citations, patent classification does capture information on the invention technical content. Rosenkopf and Nerkar (2001) and Shane (2001) propose to capture how much an invention differs from previous inventions in the field by counting the number of IPC technology classes listed in the patents cited by this patent, but in which the patent is itself not classified. However, patent (sub)classes are usually too broad to capture the detailed technical content of the invention (Arts, Hou, and Gomez 2021; Righi and Simcoe 2019; Thompson and Fox-Kean 2005). To address this issue, some other authors propose to use the text describing the invention protected by the patent. For instance, J. Kim, Y. Park, and Lee (2016) use the keywords contained in the patent document. They identify 'disruptive signals' as keywords appearing with an increasing frequency. Using a similar approach, Arts, Hou, and Gomez (2021) identify all the keywords appearing for the first time in the USPTO database. These authors provide data for all utility patents granted by the USPTO, making it easy to match with our Big Tech patents database.

Arts, Hou, and Gomez (2021)'s data will further allow us to capture the extent to which Big Tech and their targets develop disruptive innovation. It is available for all utility patents granted by the USPTO up to May 2018. By that date, 233 out of the 252 targets for which we have patent data have been granted at least one US patent. In Appendix C.1, we compare some descriptive statistics of the 252 Big Tech acquired patents portfolios with those 233 matched with Arts, Hou, and Gomez (2021)'s database.

New keywords

Arts, Hou, and Gomez (2021) build a database of the keywords extracted from the title, abstract, and claims of all the patents filed at the USPTO between 1980 and 2018. On this basis, they identify patents covering fundamentally new technologies from the number of unique keywords pairs introduced for the first time in the USPTO database that they contain.

To corroborate its validity, Arts, Hou, and Gomez (2021) show that this text-based metric outperforms the traditional measures based on patent classification (like in Rosenkopf and Nerkar 2001 and Shane 2001) and patent citations (like in Momeni and Rost 2016) in predicting whether a patent is linked to a prestigious award.[6] This result intuitively validates the capacity of new keywords to capture the impact of a novel idea.

As an extension, we also propose to more directly encompass this notion of impact in the disruption metric itself. The idea is that counting new keywords is by construction limited to how a patent stands out from *prior* knowledge, thus without reference to its impact on *future* knowledge. In order to directly integrate this notion of future impact in the disruption metric, we propose to weight the number of new keywords combinations in a patent by an index capturing the number of forward citations received by this patent (see Appendix C.2 for the construction of the metric).

[6]The following awards are considered: Nobel Prize, Lasker Award, A.M. Turing Award, National Inventor Hall of Fame, National Medal of Technology and Innovation, Benjamin Franklin Medal, and Bower Award. Additionally, the authors exploit the heterogeneity in the patent examination procedures across various patent offices, and the idea that the United States Patent and Trademark Office (USPTO) may be issuing a substantial number of weak or invalid patents. Patents granted by the USPTO but simultaneously rejected by both the European Patent Office (EPO) and the Japanese Patent Office (JPO) are assumed to lack novelty or represent only minor incremental advances over existing prior art. To assess the ability of their text-based metrics to correctly classify award patents, the authors use precision (i.e. the fraction of predicted award patents that are correctly classified), recall (i.e. the fraction of real award patents that are correctly identified) and AUC (i.e. area under the ROC-curve). We select the measure with the strongest discriminatory power to classify patents linked to prestigious awards.

3.2.2 Market environment

An important aspect of disruption relates to the market environment in which it takes place. Incumbents may not only lack the incentives to disrupt markets that they are leading (Arrow 1972, Bryan and Hovenkamp 2020a, Holmes, Levine, and Schmitz Jr 2012; Shapiro 2011), but they might also seek to hinder the disruption brought by innovative start-ups in those markets (Shapiro 2011; Federico, Morton, and Shapiro 2020).

Defining a market environment is often a delicate task. In the digital economy, where technologies are complex and not always directly tied to a specific function, circumventing the market itself can be challenging. As such, a market definition at the industry level would not be sufficiently precise. Instead, we choose to exploit a precise mapping to technological areas. Within these technological areas, different metrics can then be used to define the market environment.

In Section 3.2.2, we propose a simple metric based on the number of firms active in the same 4-digit NACE class. Of course, competition does not materialise in the same way among all the firms belonging to the same market. For a competitive analysis, we propose to use a measure of market power at the product level (Section 3.2.2). Price and cost data is rarely available at the product level, so measuring market power based on markups (i.e. ratio of the output price to its marginal cost) would not be suitable for our analysis. Instead, we choose to identify the key players of a product market based on their market shares.

Market size

As a first metric of the market environment, we use information on the "peer group size" from the Bureau Van Dijk database. A peer group is defined as companies with similar business characteristics based on the Statistical Classification of Economic Activities (NACE, at the 4-digit level). The group size, i.e the number of companies belonging to that group in the database, can then be used as a measure of market size.[7]

[7] See for instance Asdrubali and Signore (2015).

On this basis, we are able to build a measure of the size of the markets in which Big Tech acquisitions take place. Out of the 859 firms acquired by Big Tech between January 1996 and January 2021, 395 are found in the Bureau Van Dijk database. In the Appendix (Table C.1), we present the 70 groups to which Big Tech targets belong, and the number of companies in each of these groups, as of today. 43% of Big Tech targets are found within two main groups: 'Computer programming activities' (NACE 6201) and 'Other software publishing' (NACE 5829).

Market power

We aim to determine whether Big Tech's targets belong to product markets in which their acquirer is dominant.

Using market shares data from Statista,[8] we first identify the product markets where the tech giants are ranked among the key players. Statista market shares data covers the years 2019-2023, and most market data is only available for the latest year, so this measure of market power is static.

To classify Big Tech's targets in product categories, we then webscrape from crunchbase.com, for each of the identified target, its trade description and the markets in which it is active. When the Crunchbase classification is incomplete, we add additional product markets based on an online search (e.g. We add "Video Conferencing" to the list of markets associated with SKYPE.). Each target is thus associated with one or more product markets. Out of the ~2.000 sub-markets in Crunchbase data, we identify 249 different product markets to which Big Tech targets belong.

For illustrative purposes, we have classified these markets in 12 main product categories (see Appendix C.3 for the classification).[9] On Figure 3.1, we show the evolution of the number of Big Tech targets by product category. "Artificial Intelligence" and

[8] The market shares are computed by Statista based on different metrics; sales, revenue, number of users, shipment shares, app downloads, etc.

[9] In most cases, product markets associated with a same firm fall into the same product category. For firms associated with more than one product category, we manually selected the category that best fits the trade description (e.g. We selected "Application Development Software" rather than "Cloud Storage" for GITHUB).

"Entertainment" exhibit the fastest growth from 2000 onwards, overtaken by "Network Systems" since 2015.

Figure 3.1: Big Tech acquisitions over time, by product category

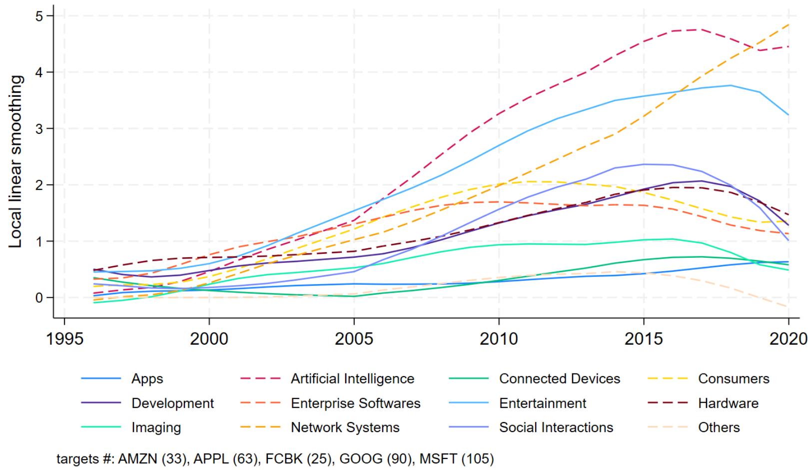

Note: The graph plots the (lowess smoothed) average number of acquisitions, separately for each of the 12 product categories in which were classified the markets in which Big Tech targets are active.

Finally, we want to identify whether a target belongs to a market in which its acquirer is dominant. To do so, we match all 249 markets from Crunchbase to markets from statista.com (see Appendix C.3 for the matching). We construct a measure of the acquirer's market dominance as a dummy capturing whether the acquirer is ranked as a key player in at least one of the markets in which its target is active:

$$Dominant_T = max\ I_{m,k}$$

where $I_{m,k} = 1$ if the market m of which the acquirer is a key player overlaps with the market k in which its target T is active.

Because our analysis will further rely on patent data, we focus on the 252 targets that have filed at least one patent application before being acquired, out of which 250 are found in Crunchbase.[10] From Table 3.1, we can see that half of the observed Big Tech targets belong to a market in which their acquirer is dominant.

Table 3.1: Big Tech acquired firms

	Count	$\text{Dominant}_T = 1$
AMZN	27	17 (63%)
APPL	52	19 (37%)
FCBK	18	6 (33%)
GOOG	67	22 (33%)
MSFT	86	62 (72%)
TOTAL	250	126 (50%)

3.3 How does Big Tech disrupt?

In this section, we want to test whether Big Tech's acquired inventions are more or less disruptive than their in-house inventions.

Let us note that our analysis is not informative regarding Big Tech's disruptive positioning with respect to a potential entrant, because we do not consider start-ups' inventions that were never acquired by Big Tech. Instead, we are comparing Big Tech's internally developed inventions and inventions developed by other firms *conditional* on being acquired by Big Tech.

To measure "disruptiveness", we use the text-based metric from Arts, Hou, and Gomez (2021): the number of unique keywords combinations contained in a patent

[10] In 2022, we webscraped market data on all 859 Big Tech targets. We then realized that some of Crunchbase market classifications were not sufficiently precise. To check for the relevance of Crunchbase classification, we aimed to further webscrape information on the target's trade description. However, in 2023, Crunchbase blocked IP addresses making requests at a high frequency, making webscraping from their website very challenging. We decided to collect the missing information manually and, to save time, we only did so for targets associated with some patent data.

document that are introduced for the first time in the USPTO database (*Disrupt*). As an extension, and to account for the patent's impact on future knowledge, we also propose to weight the number of new keywords combinations based on the number of forward citations received by the patent (*Disrupt**).

Our sample thus covers all the patents granted to the 5 tech giants and their 233 targets with some patented inventions matched with textual data. We observe that Big Tech tends to publish proportionally more of very little disruptive patents (with 0 new keywords combination, see Figure C.2 in the Appendix) than their targets. This is likely because the tech giants can more easily afford to make bad investments from time to time. To account for this, we will test for the robustness of our results to the exclusion of patents with a disruption value of zero.

When aggregating across firms (i.e. patent portfolios), we can see in the Appendix that Big Tech is, on average, less disruptive than its targets (see Tables C.5 and C.6). This result also holds when we exclude patents with a disruption value of zero (see Tables C.7 and C.8).

3.3.1 Comparing internally developed and acquired inventions

A difference in the average disruption metrics between internally developed and acquired patent portfolios might not be fully attributable to different disruption levels. Our goal is to compare inventions that are developed internally and those that are acquired by Big Tech while keeping constant other determinants of our disruption metric. To do so, we construct the analysis at the patent level, which allows to capture more variation in other potential endogenous determinants of the disruption metric.

Relevant patents

A first reason to refine the analysis is that not all patents in a portfolio are relevant for assessing the disruptiveness of technologies developed by a firm. This is because most of the information is contained into patents that are very disruptive, but these might be hidden in a mass of very little disruptive patents. In the literature, this is sometimes dealt with by selecting the top percentiles of the distribution of patents.[11] However, in our analysis, we are comparing patent portfolios that are very different in sizes (with around 15.500 patents for Big Tech vs. 16 patents for their targets). So selecting, for instance, the top 10% patents at Microsoft is very different from selecting the top 10% patents at Skype. Instead, we propose to adapt the selected values based on the company size.

As a proxy for the company size, and since we are mainly interested in the innovation potential of the company, we use the number of inventors who have published some patent for the company. This variable is constructed based on the number of unique inventor id(s) in the firm's patent portfolio. Next, we sort all the patents published by this firm based on their disruption indices. Our sample is then restricted to all the patents that are ranked in the top x, where x is the number of inventors at the firm. When the number of patents in the firm's portfolio is lower than the number of inventors, we just keep all the patents.

Other determinants of the disruption metric

By construction, our disruption metrics are likely to be systematically higher for patents containing more text, and hence more keywords. To correct for this first potential source of endogeneity, we include two proxies for the patent text length: the number of claims (*ClaimsNbr*) and the number of distinct 4-digit IPC technology subclasses (*PatentScope*)[12] contained in the firm's patent portfolio.

Next, to control for potential scale effects, we include the average number of countries in which the inventions patented by the firm are protected (*FamilySize*), and dummies capturing the industries to which the patent belongs (*CPC*). While GAFA platforms assume similar roles in online activities, Microsoft is sometimes considered separately (Galloway 2018, Simon and Joel 2011). For this reason, we also include a dummy variable taking the value of 1 for Microsoft's - published or acquired - top patents (*MSFT*).

[11] See for instance Morzenti 2022.

[12] Inventions at the intersection of several technology classes (likely common for big companies like the tech giants) might contain more text.

Finally, we include dummies for the patent grant year (G). This controls for the fact that Big Tech's patents are likely to have been, on average, granted earlier in time than its targets', while our measures of disruption may be time-sensitive.

Model

To ensure that the differences in disruption between Big Tech's top patents ($Acquired = 0$) and its targets' ($Acquired = 1$) can be attributed to the identity of the firm behind the protected inventions, we need to keep constant the identified potential endogenous determinants of disruption:

$$\begin{aligned} &Disrupt_p = \\ &\qquad f(Acquired_p, MSFT_p, PatentScope_p, ClaimsNbr_p, FamilySize_p, CPC_p, G_p) + \varepsilon_p \end{aligned} \tag{3.1}$$

We first define the disruptiveness of a patent p ($Disrupt_p$) as the number of new keywords combinations contained in the patent. This variable only takes non-negative values, and its distribution is skewed to the left (see Appendix C.4); most observations are associated with very few new keywords combinations. To correct for this over-dispersion observed in the distribution of the dependent variable, we fit a Negative Binomial regression model, which allows for a variance greater than the mean (Wooldridge 2010, pp.657-659):

$$\begin{aligned} Var(Disrupt_p) = \; & E\left(\beta_0 + \beta_1 Acquired_p + \beta_2 MSFT_p + \beta_3 Acquired_p MSFT_p + z_p\gamma + \eta G_p\right) \\ & + \eta^2 \left[E\left(\beta_0 + \beta_1 Acquired_p + \beta_2 MSFT_p + \beta_3 Acquired_p MSFT_p + z_p\gamma + \eta G_p\right)\right]^2 \end{aligned}$$

where the vector z contains the control variables ($PatentScope$, $ClaimsNbr$, $FamilySize$ and CPC dummies), and η^2 is the variance of the error term.

Model (3.1) is estimated by Maximum Likelihood. In the log-likelihood function, we use the probability mass function of the negative binomial distribution. The obtained coefficients estimates are presented in Table 3.2, and in the Appendix for the alternative measures of disruption.

Table 3.2: Disruption of Internally developed vs Acquired top patents

	(1)	(2)	(3)	(4)	(5)
Disrupt					
Acquired=1	0.538***	0.526***	0.524***	0.564***	0.526***
	(14.84)	(14.43)	(14.37)	(15.62)	(14.43)
MSFT=1	0.434***	0.435***	0.436***	0.552***	0.553***
	(31.03)	(31.06)	(31.11)	(38.79)	(38.35)
Acquired=1 × MSFT=1	-0.481***	-0.473***	-0.479***	-0.597***	-0.541***
	(-7.36)	(-7.22)	(-7.29)	(-9.17)	(-8.29)
Patent Scope		0.021***	0.017***	0.014**	0.014**
		(3.66)	(2.90)	(2.31)	(2.20)
Family Size			0.010***	0.001	0.002
			(4.72)	(0.38)	(0.80)
Number of Claims				0.028***	0.028***
				(39.70)	(39.57)
Constant	3.367**	3.346**	3.340**	3.041*	3.159*
	(1.99)	(1.98)	(1.98)	(1.82)	(1.90)
Observations	76764	76446	76446	76444	76440
Pseudo R^2	0.012	0.012	0.012	0.015	0.016
Year dummies	Yes	Yes	Yes	Yes	Yes
CPC dummies	No	No	No	No	Yes

t statistics in parentheses

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

We find that the parameters associated with acquired (vs internally developed) top patents are positive for all tech giants except for Microsoft: $\hat{\beta}_1 = .526^{***}$ for GAFA, and $\hat{\beta}_1 + \hat{\beta}_3 = -.015$ for Microsoft. This means that, after controlling for other factors that may impact our first measure of disruption, acquired inventions are, on average, more disruptive than internally developed inventions (or, in the case of Microsoft, not significantly different). This result is robust to weighting the number of new keywords combinations based on the number of forward citations: $\hat{\beta}_1 = .790^{***}$ and $\hat{\beta}_1 + \hat{\beta}_3 = .174^{**}$ (see Table C.9 in the Appendix), and excluding patents with zero disruptive value: $\hat{\beta}_1 = .486^{***}$ and $\hat{\beta}_1 + \hat{\beta}_3 = .020$ (see Table C.10).

3.4 Where does Big Tech disrupt?

To answer our first research question, we compared the top patents of Big Tech platforms and Big Tech-acquired firms, and we concluded that Big Tech's inventions are on average less disruptive than inventions developed by its targets. However, this difference does not necessarily stem from Big Tech's failure to disrupt. Instead, Big Tech might see the acquisition of disruptive start-ups as a substitute to bringing in-house disruptive innovation in technological markets where it is less strong. Conversely, the acquisition of disruptive start-ups can be viewed as a tactic to stifle disruption in markets that the tech giants are leading (Lemley and Wansley 2024; Shapiro 2011). If such strategic considerations were to underline Big Tech's disruption effort, we should observe different disruption patterns depending on the market environment in which the innovative activities take place.

3.4.1 Technology adoption

To trace disruptive technologies later adopted by their acquirer, we compare the text of patents in the target firms' portfolio and in the acquirer's portfolio.

Based on Arts, Hou, and Gomez (2021), we identified disruptive technologies from combinations of keywords introduced for the first time in history by granted US utility patents. We now postulate that, when the acquirer is adopting a disruptive technology on which the target was working, the associated new keywords combinations will appear in the acquirer's own patent portfolio: $Adopt_{a,t}$ takes the value of the number of new keywords combination(s) from a given patent t of the target T's portfolio found in another patent a of the acquirer's portfolio. We construct the variable $Adopt_T$ as the total number of new combinations of keywords further reused by the acquirer:

$$Adopt_T \equiv \sum_a \sum_t Adopt_{a,t}$$

Since technology adoption is measured up until May 2018 (end of Arts, Hou, and Gomez (2021)'s study period), we restrict our analysis to acquisitions that have been undertaken up until that date: these represent 196 out of the 225 targets for which textual data is available and that have patented some disruptive technology (i.e. that are associated with at least one new keywords combination).

As can be seen from Table 3.3, for about 3/4 of Big Tech acquisitions, a keywords combination introduced for the first time in history by a target patent is further found in its acquirer's portfolio. So it seems that a large majority of Big Tech acquisitions involve the adoption of the target's disruptive technology.

Table 3.3: Big Tech acquired patents portfolios

	Count	$Adopt_T > 1$
AMZN	19	14 (74%)
APPL	41	25 (61%)
FCBK	11	8 (73%)
GOOG	57	40 (70%)
MSFT	68	54 (79%)
TOTAL	196	141 (72%)

This count includes Big Tech patents portfolios acquired by May 2018 with textual data and at least one new keywords combination.

To compare the variable $Adopt_T$ across patent portfolios, we must account for a size effect; for the fact that targets with a higher number of disruptive technologies are more likely to have some of them later adopted by their acquirer, and that more innovative acquirers are more likely to adopt the acquired technology. This will further be captured by the number of acquirer/target units combinations:

$$Combi_T = NewKeywords_T \cdot PatentsNbr_A$$

where $NewKeywords_T$ is the number of new combinations of keywords in the target's patent portfolio, and $PatentsNbr_A$ is the number of patents in the acquirer's portfolio.[13]

[13] This term also captures variations in the length of the period during which the acquirer could potentially

3.4.2 Disruptive technology adoption across market environments

To test whether our results vary depending on the market environment in which these innovative activities take place, we focus on two features of the market: the market size and the acquirer's market power. Because the market environment pertains to interactions between firms, this market data is collected at the firm level, so we run the analysis at the target level.

Model

We assume a Negative Binomial distribution of the technology adoption variable to account for the over-dispersion observed in the data; a majority of disruptive technologies are not adopted (i.e. new combinations of keywords never reused) by the acquirer, but a few targets are associated with many new combinations of keywords further reused by the acquirer. On this basis, we can test whether the disruptive technology adoption depends on the market environment in which the acquisition takes place:

$$Adopt_T = f(X_T, Combi_T, Z_T) + v_T \tag{4}$$

The regressor of interest (X_T) either captures the size of the market in which the acquisition takes place ($PeerSize_T$), or the acquirer's market power ($Dominant_T \in \{0, 1\}$). In addition to the size effect, captured by $Combi_T$, we also control for some target's characteristics in Z_T: the year in which the target filed a patent for the first time ($FirstFiling_T$), whether the target has filed some patent originating from the United States ($US_T \in \{0, 1\}$), and the year in which the target is acquired ($AcquiYear_T$).

We thus want to test whether Big Tech's adoption of disruptive technologies from acquired start-ups varies with the number of firms/competitive pressure in the product markets to which the start-ups belong. Let us recall that Statista market power data covers the period 2019-2023 and Orbis market size data has been collected in 2024, while technology adoption is measured up until May 2018. So what we are testing is whether the tech giants have adopted more disruptive technologies (historically) in technology

adopt the disruptive technology, since $PatentsNbr_A$ is higher for earlier acquisitions.

fields in which competition is weaker (today).

The coefficients estimates, obtained by Maximum Likelihood with the probability mass function of the negative binomial distribution, are presented in Tables 3.4 and 3.5. In the last column, we present the parameters estimated based on acquisitions up to 2016 (thus excluding 2017 and 2018). This restriction aims to test whether our results are robust to excluding acquired technologies that have been in their acquirer's portfolio for less than 1.5 year, and thus might not have been adopted *yet*.

Table 3.4: Market size and Disruptive technology adoption

	(1) All	(2) All	(3) All	(4) Excl. >=2017	(5) Excl. >=2017
Adopt					
PeerSize	-0.000	-0.000	-0.000	-0.000	0.000
	(-0.61)	(-0.73)	(-0.24)	(-0.07)	(0.01)
Combi	0.000^{**}	0.000^{***}	0.000^{***}	0.000^{**}	0.000^{***}
	(2.27)	(2.75)	(3.21)	(2.18)	(3.18)
US		1.054^{*}	0.474	0.160	0.100
		(1.78)	(1.03)	(0.26)	(0.21)
FirstFiling			-0.066		-0.049
			(-1.64)		(-1.27)
Constant	-21.950	-22.892	135.907^{*}	-19.232	102.424
	(-0.00)	(-0.00)	(1.69)	(-0.00)	(1.32)
Observations	124	124	124	106	106
Pseudo R^2	0.041	0.044	0.025	0.031	0.025
AcquiYear dummies	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.5: Acquirer's market power and Disruptive technology adoption

	(1) All	(2) All	(3) All	(4) Excl. >=2017	(5) Excl. >=2017
Adopt					
Dominant	-1.128***	-1.095***	-1.020**	-1.021**	-0.914**
	(-2.79)	(-2.73)	(-2.53)	(-2.50)	(-2.26)
Combi	0.000***	0.000***	0.000***	0.000***	0.000***
	(3.92)	(4.06)	(3.53)	(3.95)	(3.41)
US		0.403	0.496	0.156	0.344
		(1.07)	(1.26)	(0.40)	(0.82)
FirstFiling			-0.053		-0.080
			(-0.93)		(-1.32)
Constant	9.780***	9.339***	114.654	9.526***	169.093
	(6.69)	(6.16)	(1.01)	(6.58)	(1.40)
Observations	191	191	191	169	169
Pseudo R^2	0.052	0.053	0.053	0.049	0.050
AcquiYear dummies	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

While the disruptive technology adoption does not seem to vary across market sizes (see Table 3.4), we do observe a negative relationship with the acquirer's market dominance (see Table 3.5): Big Tech has adopted (historically), on average, 60% to 64% (= $1 - e^{\hat{\beta}}$) fewer disruptive technologies in markets in which it is dominant (today). In other words, the acquirer's market power negatively relates with the adoption of disruptive technologies from acquired start-ups.

3.5 Conclusion

This paper sheds light on the adoption of disruptive innovation in digital markets. We study the innovative activities of the tech giants, both in-house and through the acquisition of disruptive start-ups.

On the one hand, large Big Tech platforms are often considered insufficiently responsive to develop in-house disruptive innovation, so they compensate by acquiring smaller, more agile start-ups (Bresnahan, Greenstein, and R. M. Henderson 2011; R. M. Henderson and Clark 1990). On the other hand, these tech giants might not face the right incentives to disrupt markets that they are leading (Bryan and Hovenkamp 2020a; Holmes, Levine, and Schmitz Jr 2012) and, instead of adopting the acquired technology, they might aim to shelve it to avoid directing users away from their existing services (Lemley and Wansley 2024; Shapiro 2011). In line with this existing literature, our analysis reveals nuanced insights into the relationship between Big Tech platforms and disruption, with findings that vary with the origin of the invention and the competitive context in which it is developed.

Our results first suggest that Big Tech's acquired inventions are more disruptive than those that it develops internally. However, this difference does not necessarily stem from Big Tech's failure to disrupt, as the acquisition of disruptive start-ups might be used as a substitute to bringing in-house disruptive innovation in technological markets where it is less strong. In line with this hypothesis, we find that disruptive technologies adopted through acquisition predominantly emerged from markets where the tech giants hold a weaker position.

These results suggest that inventions developed by the tech giants tend to be relatively little disruptive and that, to remain on the technology edge in markets in which they hold a weaker position, these leading firms often resort to acquiring disruptive start-ups. As a potential improvement to this analysis, future research could focus on developing a dynamic measure of Big Tech's market power across its various technological domains. This would provide a clearer understanding of Big Tech's ex-ante motivations for adopting disruptive technologies both within and outside of dominated markets. In addition, further research could offer valuable insights into the strategies employed by the other main players of the digital sector, e.g. Tesla, Samsung, Oracle, Alibaba, Adobe, IBM.

Appendix A

A.1 Timing of the patenting process

Before a patent is granted, it must be filed and published. The legal requirement for the patent office to publish a patent application is 18 months from the filing. This 18-month limit is respected for 95% of all US patent applications (Tegernsee 2012). Earlier publication is often observed: half of US patent applications are published within 9 months after they were filed (Martin 2015).

Publication means that the content of a patent application is known to the public; that is becomes "prior art". However, it does not necessarily mean that the application will result in a (granted) patent, which grants to the applicant the exclusive rights over the use and sale of the invention. On average, US patents are granted within 32 months of their filing date (as computed based on the 'grant lag' from the OECD Patent Quality Indicators database, July 2021).

A.2 Focus on patent-protected technologies

Figure A.1: Big Tech targets with and without patents, by funding amounts

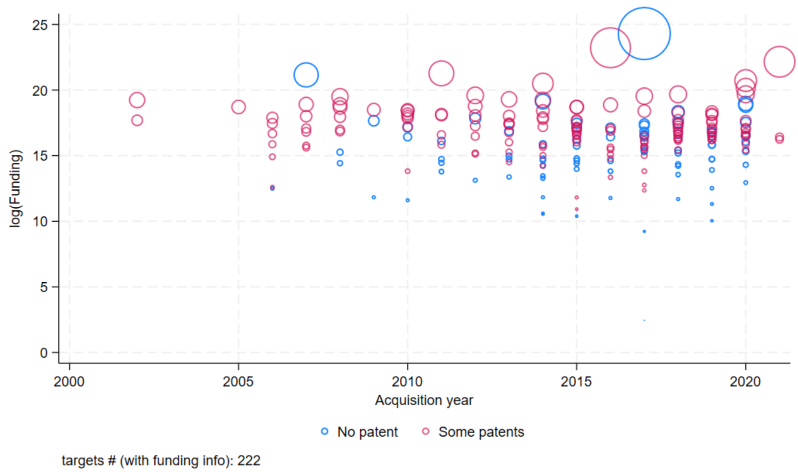

Notes: For readability, the funding amount variable is log-transformed and the sizes of the circles are a function of its values. The two largest blue circles are associated with AQUANTIVE (acquired by Microsoft in May 2007) and WHOLE FOODS MARKET (acquired by Amazon in June 2017).

A.3 Citations data to capture technology developments

In our study, we use patent citations as a proxy for the innovation effort in a given technology field. Because all previous knowledge used in an innovation has to be cited in the patents protecting this innovation, if a technology stops being developed, one should observe fewer citations to the patents protecting this technology. On the contrary, a technology that is further developed will be cited in many subsequent patents. Information about patents citations is therefore very useful to study Big Tech's acquisition strategies, because it allows to infer the use that is made of an acquired technology in subsequent innovation. More specifically, we can capture the improvements that are made by an acquirer to an acquired technology based on the number of acquirer's citations to the patents protecting that technology.

Of course, using patent data to identify changes in the acquired technology development suffers from an important limitation; it only accounts for patent-protected technologies. Some innovations might not have been patented, because they are simply not patentable or due to high costs of patenting (e.g. hiring patent specialists to prepare the application, paying the filing administrative costs and the renewal fees).

Information on the number of forward citations made to a given patent also suffers from some biases. Companies might have strategic reasons not to cite a patent. For instance, fewer citations would be made by firms aiming to gather patents for defensive or cross-licensing purposes (Abrams, Akcigit, and Grennan 2013; Jaffe, Trajtenberg, and R. Henderson 1993; Lampe 2012). This should not be a problem in our analysis as we do not only consider citations made by the applicant, but also those added by the examiner. Citations data might also be noisy (Gambardella, Harhoff, and Verspagen 2008) due to differences between applicants (Rysman and Simcoe 2008; B. N. Sampat 2010) and across industries (Lerner, Sorensen, and Strömberg 2011; Rysman and Simcoe 2008). For our analysis, we focus on the digital sector, so cross-industry heterogeneity should not affect our results. Our study of the evolution of citations made by Big Tech is also little affected, since we consider the same five applicants over time. Another potential source of bias is that the citations count might include irrelevant references as patent applicants have an incentive to cite as many references as possible; if a reference the applicant knew about is forgotten, a court may rule the patent to be unenforceable in infringement proceedings (Allison and Lemley 1998; Kuhn, Younge, and Marco 2020). But the resulting measurement error has been shown to be mainly problematic for the study of citation patterns over time (Kuhn, Younge, and Marco 2020; Marco 2007), so this can be accounted for in our analysis by controlling for the date at which a given citation is observed. Finally, front page citations, based on which the USPTO forward citations that we are using are constructed, include some, but not all, of the citations contained in the patent text. It has been shown that in-text citations might be a better measure of the development of the invention (Bryan, Ozcan, and B. Sampat 2020). However, these in-text citations are seldom used by innovation researchers.[1] This is because there is no standardized format for in-text citations, so they are difficult to extract (Narin and Noma 1985).[2]

[1] The only study of in-text references that we are aware of, carried out by Bryan and Ozcan (2021), uses a matching method to search the text of all patent applications for references to articles in medical journals.

[2] We are not aware of any large-scale database on in-text patent citations.

A.4 Observational cut in the citations database

Figure A.2: Distribution of filing dates for all citing patents (Density)

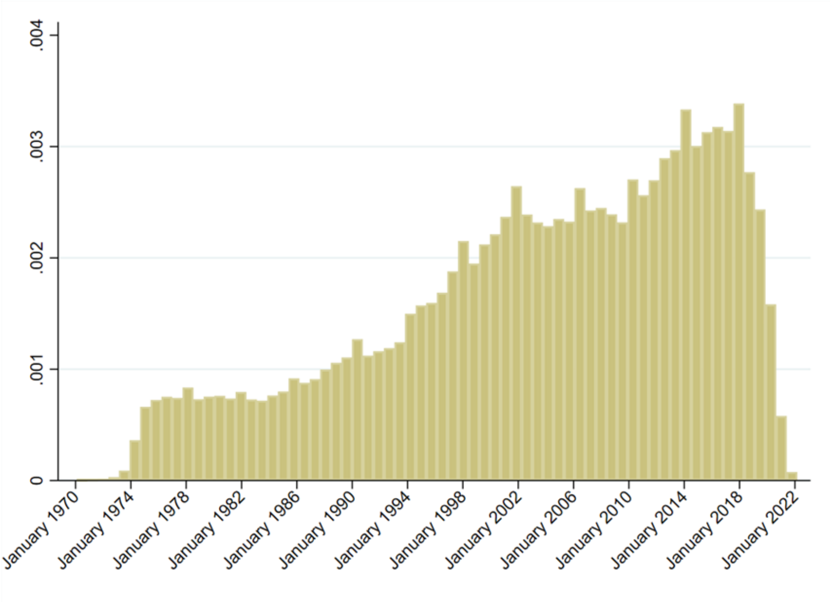

Notes: For clarity, the filing dates before 1970m1 (2% of the sample) are not represented.

We observe a drop in citations after January 2018 because citing patents have not been granted yet. For this reason, we end up our sample in June 2017.

A.5 Negative Binomial distribution of the citations count

Figure A.3: Distribution of the count of Big Tech citations (Percent)

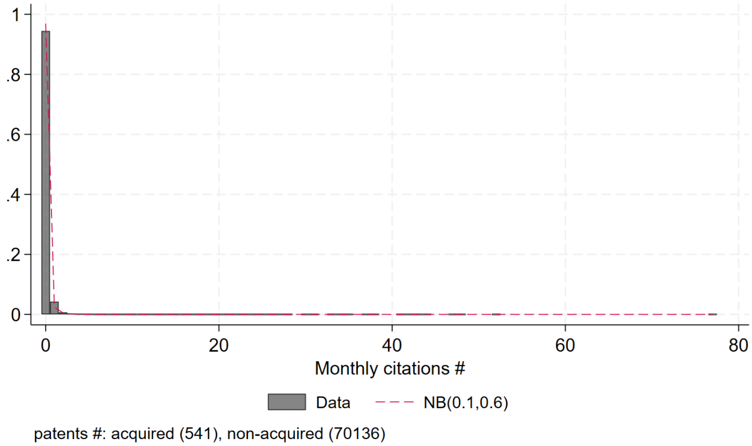

Notes: This figure shows an histogram of the number of citations received by a given patent in a given month, overlaid with a negative binomial density with the same parameters.

A.6 Inverse probability weighting

In order to make acquired and non-acquired patents comparable in all respects except for their acquisition status, as if acquisition had been fully randomized, we use *propensity scores*. Propensity scores can be seen as the channel through which a patent's characteristics affect its acquisition status and hence create endogeneity in the relation between the treatment (the acquisition status) and the outcome (forward citations). Because most determinants of both a patent's acquisition status and the citations it receives are unobserved, they will be controlled for by using the pre-treatement outcomes (i.e. pre-acquisition patent citations).

We first estimate a discrete choice Probit model of the probability for a patent p to have been acquired $P(A_p = 1)$ with, as regressors, the citations this patent receives pre-acquisition, both in levels ($Cit_{p,Pre}$) and in growth rates ($CitGR_{p,Pre}$):[3]

$$P\left(A_p = 1|CitGR_{p,Pre}, Cit_{p,Pre}\right) = \Phi(\alpha + \beta^1 CitGR_{p,Pre} + \beta^2 Cit_{p,Pre}) \tag{A.1}$$

where $CitGR_{p,Pre}$ captures the growth rate in the number of citations between the first and the last periods pre-acquisition ($t = -1$ and $t = -3$), $Cit_{p,Pre}$ captures the number of citations in $t = -2$, and Φ is the cumulative density function of the standard normal distribution.

We then use the predicted values from the function to generate, for each observation, the propensity scores (P_p), which ensure that patents with the same pre-acquisition citations have a positive probability of being both acquired and non-acquired.

Next, to disentangle the effect of acquisition from the effect of potential confounding factors, we need to close the propensity scores channel through which these confounding factors affect a patent's acquisition status. This can be done by using the propensity scores to conduct *inverse probability weighting* (King and Nielsen 2019). The first step of this procedure consists in "trimming" non-acquired patents outside of the acquired patents' propensity score range. This limits the data to the range of "common support", i.e. to non-acquired patents that are sufficiently comparable to acquired patents. Second, we need to weight each acquired patent by the inverse of the probability that it was acquired ($1/P_p$), and each non-acquired patent by the inverse of the probability that it was not acquired ($1/(1 - P_p)$). By weighting patents by the inverse of the probability of what they actually are, we make the treated and control groups more similar. Acquired patents that get the biggest weights are the ones that are most like non-acquired patents; acquired patents who were least likely to have been acquired. Inversely, non-acquired patents with the biggest weights are the ones most like acquired patents; non-acquired patents who were most likely to have been acquired (Huntington-Klein 2021). In turn, we obtain a sample of patents in which individual heterogeneity has been averaged across the treatment and control groups.

[3] While the large sample sizes would allow to use additional regressors, like the IPC technology class, doing so would not lead to pre-acquisition parallel trends of acquired and non-acquired patents' citations in the estimation of model 1.2.

To ensure that this re-weighting will properly take out the effect of endogenous covariates on the acquisition status, we must test for "balance". In our case, balance means that, after weighting, there are no more meaningful differences between acquired and non-acquired patents in pre-acquisition citations. This ensures that the inverse probability weighting is appropriate to close the propensity scores channel through which confounding factors affect a patent's acquisition status, i.e. that acquired and non-acquired patents become similar in all aspects except for their acquisition status. A common way of checking for balance is to test for the difference of means between the control and the treated groups. Table A.1 presents the results of this test before and after applying the inverse probability weighting. We observe that the differences in citations means before (simulated) acquisition between acquired and non-acquired patents are reduced (.062 in the raw sample, .059 in the new trimmed and weighted sample). This exercise illustrates how dropping observations outside the range of common support and weighing observations based on their inverse probabilities allows a better comparison of the two patent groups post-acquisition. However, since we are interested in the evolution of citations around acquisition time, the most important condition for a meaningful comparison of the two groups is the pre-acquisition parallel trends in the estimated DIS (see Figure 1.3).

Table A.1: Balance tables

Raw sample (before trimming and weighting)

Variable	(1) Not acquired	(2) Acquired	(3) Acquired vs Not
$\overline{Cit_{Pre}}$	0.232	0.294	0.062
	(0.869)	(0.806)	(0.037)
Observations	77,522	541	78,063

Working sample (after trimming and weighting)

Variable	(1) Not acquired	(2) Acquired	(3) Acquired vs Not
$\overline{Cit_{Pre}}$	0.217	0.276	0.059
	(0.758)	(0.735)	(0.004)
Observations	77,359	541	77,900

These tables present the results of the balancing test for the inverse probability weighting. In the first and second columns, we show the means and the standard deviations of the pre-acquisition citations, for control observations (non-acquired patents) and treated observations (acquired patents) respectively. In the third column, we regress those pre-acquisition citations on the observation's treatment value (acquired or not) to compute the differences of means and the associated standard errors.

A.7 Sharp event study - Identification

In a paper studying the impact of having children on the gender wage gap, Kleven, Landais, and Søgaard (2019) exploit the sharp breaks in career trajectories occurring just after the birth of a child. We present below the conceptual framework set out by these authors, adapted to our research question.

The number of citations made at time t by an acquirer to some acquired patent p is defined as a function of variables in $x_{p,t}$ responding to an acquisition event (such as the type of portfolio in which the newly acquired patent is integrated), and variables in $z_{p,t}$ that do not depend on acquisition (such as the age of the patent, its quality, characteristics of the publishing company, etc.):

$$Cit_{p,t} = f(J'\tau_j, x(J', z_{p,t})\tau_x, z_{p,t}\tau_z) + \varepsilon_{p,t} \tag{A.2}$$

where $J' = \sum_{j \neq 0} Ij = t$ is a vector indicating the time at which the citation is observed with respect to the time of acquisition. In this framework, citations may respond directly to acquisition conditional on $x_{p,t}$, and indirectly through $x_{p,t}$ (e.g. the impact of complementarities/substitutions with other patents from the new portfolio).

For changes in the number of citations to correctly identify the post-acquisition impacts, the first condition is that "the event" should not determined by the outcome variable. In our case, this implies that, conditional on the set of underlying determinants $z_{p,t}$, acquisition is exogenous to the outcome variable $Cit_{p,t}$. To set up the additional necessary conditions under which we can identify the effect of acquisition, we must distinguish between the short-run and the long-run.

Our identification strategy of the short-run effect of acquisition relies on one additional assumption: the event should generate sharp changes in the outcomes that are orthogonal to unobserved outcome determinants. This 'smoothness assumption' is needed because, when we shock J, we get a response in the number of citations that is captured by both τ_j and τ_x. But τ_x does not only respond to the event time; it also captures the effect of changes in the variables in $z_{p,t}$, which could happen at the same time as acquisition. However, if we assume that citations would evolve smoothly absent acquisition, the short-run effect of acquisition conditional on $z_{p,t+}$ can be identified from the change in the number of citations when going from the acquisition time (t_0) to an event time just after (t_+):

$$E\left[Cit_{p,t+} - Cit_{p,t0}\right] = E[f(1, x(1, z_{p,t+}), z_{p,t+}] - E[f(0, x(0, z_{p,t0}), z_{p,t0}] \qquad \text{(A.3)}$$

where the smoothness of the average citations path absent acquisition would imply that $E[F(0, x(0, z_{p,t+}), z_{p,t+}] \approx E[F(0, x(0, z_{p,t0}), z_{p,t0}]$. The short-run impact of acquisition is therefore identified from the sharp changes in citations immediately following acquisition rather than from the smooth trends in citations. The graphical evidence presented on Figure 1.1 lends support to the suitability of this conceptual framework for our analysis, as the sharp breaks in citations trajectories occurs *just after* acquisition.

The long-run impact is obtained by considering an event time t_{++} long after the acquisition time:

$$E\left[Cit_{p,t++} - Cit_{p,t0}\right] = E[f(T, x(T, z_{p,t++}), z_{p,t++}] - E[f(0, x(0, zp, t0), zp, t0] \qquad \text{(A.4)}$$

The differences between this impact measure and equation A.3 is that the smoothness assumption is no longer sufficient for identification as we can still have large changes in citations determinants (other than acquisition) over a long event time window. Indirectly controlling for $z_{p,t}$ with age and date dummies, as we do in model 1.1, may partially solve this problem. But we cannot claim that we have controlled for all elements of $z_{p,t}$, so the event study estimates representing the change in the number of citations compared to the time of acquisition (θ^1 in model 1.1) might not properly capture the long-run impact of acquisition. We therefore propose with model 1.2 a second solution to capture long-term effects of acquisition, by using a control group to account for the citations trend absent acquisition.

A.8 Robustness checks

Additional regressors

We rewrite model 1.1 to control for additional determinants of the number of citations received by a patent:

$$Cit_{p,j,t,d} = f(J'\theta^6, age_{p,d}\beta^6, M'\gamma^6, firm_j\xi^4, Z_p\nu^1) + \varepsilon^4_{p,j,t,d} \tag{A.5}$$

where Z_p contains the additional regressor(s).

Microsoft First, we would like to control for the potential effect specific to those patents acquired by Microsoft. While GAFA platforms assume similar roles in online activities, Microsoft is sometimes considered separately (Galloway 2018, Simon and Joel 2011). Furthermore, Microsoft is the biggest acquirer in our sample (see Table 1.2). In this case, Z_p takes the value 1 if patent p was acquired by Microsoft, 0 if it was acquired by Google, Apple, Facebook or Amazon. The estimated incident rate ratios are presented on Figure A.4 (a).

More citations determinants Second, Z_p is defined such as to contain two additional citations determinants: the technology field to which patent p belongs, and the origin of its publishing company. Of all the patents published by Big Tech, 57% and 32% contain at least one reference to a technology field classified in the CPC sections "Physics" and "Electricity", respectively. The other CPC fields are barely represented in Big Tech patent portfolios, with frequencies going from 0% to 2%. We include in Z_p two dummy variables, one for "Physics" and one for "Electricity", indicating whether patent p is associated with that technology field. In addition, we include an indicator variable capturing whether the company that published the patent was located in the US (77% of our working sample), in the EU (13%) or in the Middle East (10%). The estimated incident rate ratios are presented on Figure A.4 (b).

Figure A.4: Big Tech citations to acquired patents relative to acquisition, more controls

(a) Allowing for a Microsoft-effect (b) Controlling for CPC fields and publisher's origin

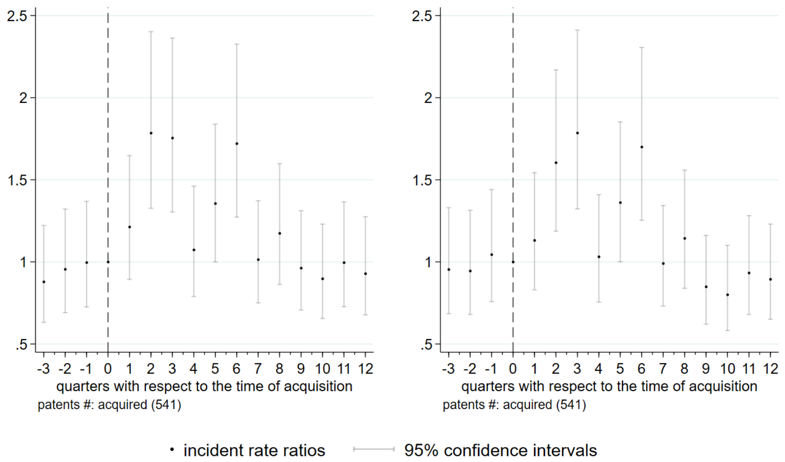

Notes: The graph shows the incident rate ratios for acquired patents: $e^{\hat{\theta}_t^6}$ from model A.5. These coefficients are estimated on a balanced sample of patents in a 4 year-window around acquisition.

The inclusion of these additional regressors seem to have little impact on our results, as the estimates presented on Figure A.4 appear to be very similar to our baseline results presented on Figure 1.2.

Alternative study periods

Extending the pre-treatment period On Figures A.5 and A.6, we replicate the results from Models 1.1 and 1.2 for a pre-treatment period of 15 months (instead of 9 months). The coefficients estimates follow very similar trajectories to those in our baseline results.

We note significant incident rate ratios in the earliest quarters before acquisition ($t = -5$ and $t = -4$). We argue that anything happening one year or more before acquisition is unlikely to be relevant to the acquisition event, and that the 'no-acquisition' counterfactual can thus be specified based on the last portion of the pre-intervention period ($t = -3$ to $t = -1$).

Figure A.5: Big Tech citations to acquired patents relative to acquisition, longer pre-treat

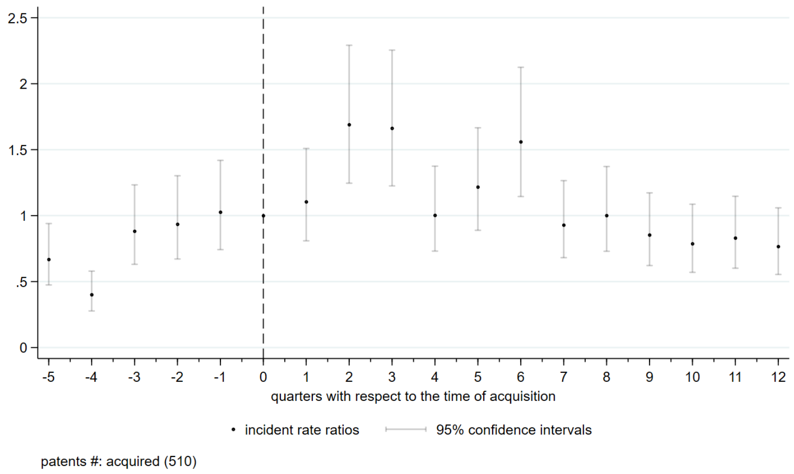

Notes: The graph shows the incident rate ratios for acquired patents: $e^{\hat{\theta}_t^1}$ from model 1.1. These coefficients are estimated on a balanced sample of patents in a 4.5 year-window around acquisition.

Figure A.6: Big Tech citations to acquired patents w.r.t. non-acquired patents, relative to the (simulated) acquisition announcement, longer pre-treat

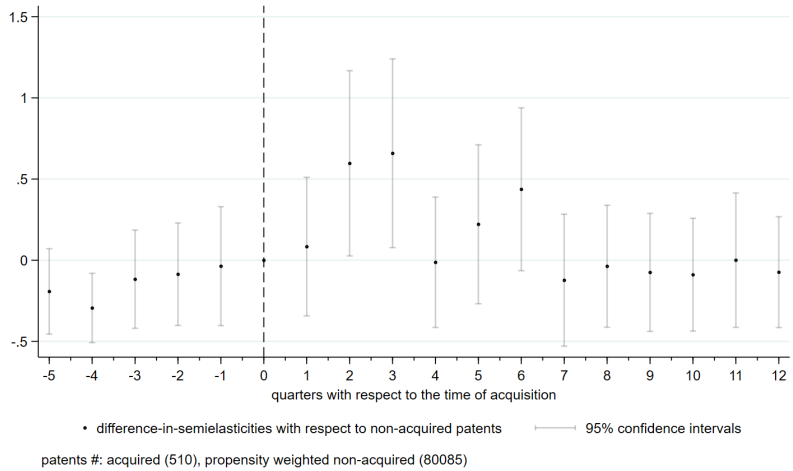

Notes: The graph shows the DIS between acquired and non-acquired patents: $e^{(\widehat{\theta}_t^2+\widehat{\alpha}_t^1)} - e^{(\widehat{\theta}_t^2)}$ from model 1.2. These coefficients are estimated on a balanced sample of patents in a 4.5 year-window around (simulated) acquisition.

Reducing the study period Next, we reduce our study period to 2 (instead of 4) years around acquisition; 1 quarter before acquisition, 6 quarters after. This allows to consider some Big Tech-acquired firms that were not included in our baseline sample: i. those that only started patenting shortly before being acquired, ii. those acquired between May 2014 and January 2016.[4] On Figure A.7, we observe that the evolution of citations just after acquisition follows a very similar trend to the baseline sample: citations increase significantly after acquisition. Since the observation period is reduced, we do not capture the slow down in citations after 1.5 year.

[4]Because we end our study period in June 2017 to avoid biases in the citations count, restricting our baseline sample to patents observed up to 3 years after acquisition meant that we could only use acquisitions undertaken until May 2014 (58% of all 855 Big Tech acquisitions from Table 1.1). By including patents observed up to 1.5 year (instead of 3 years) after acquisition, we capture acquisitions until December 2015 (72% of all 855 Big Tech acquisitions).

Figure A.7: Big Tech citations to acquired patents relative to acquisition, reduced period

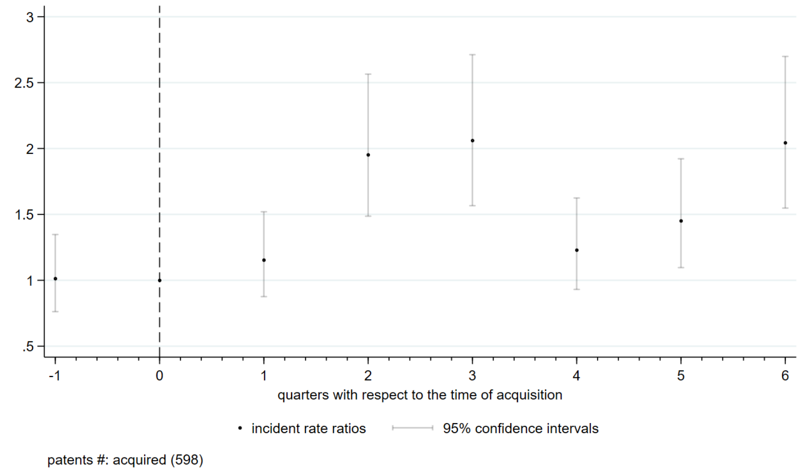

Notes: The graph shows the incident rate ratios for acquired patents: $e^{\hat{\theta}_t^1}$ from model 1.1. These coefficients are estimated on a balanced sample of patents in a 2 year-window around acquisition.

Including Motorola As an additional check, we include *Motorola Mobility*, acquired by Google in August 2011 and later (January 2014) sold to Lenovo. Motorola was not included in our baseline sample because its acquisition status changes during the study period. Its patents belong to Google for only 29 months after acquisition, while our study period covers three years after acquisition. However, since Motorola has a very large patent portfolio, owning 1,080 patents at acquisition among which 125 are cited by Google by June 2017, we propose an alternative study period that allows to include it. We can see from Figure A.8 that our baseline results are robust to this alternative specification.

Figure A.8: Big Tech citations to acquired patents relative to acquisition, with Motorola

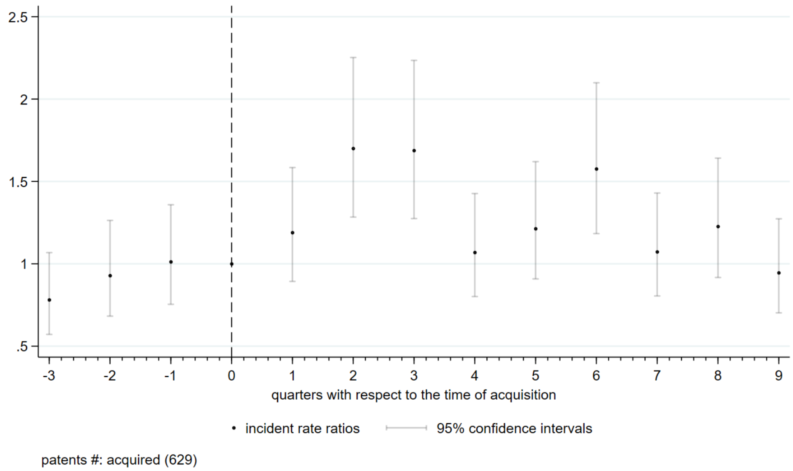

Notes: The graph shows the incident rate ratios for acquired patents: $e^{\hat{\theta}_t^1}$ from model 1.1. These coefficients are estimated on a balanced sample of patents in a 3 year-window around acquisition.

Heterogeneity-robust treatment effects

To allow for variations in the effect of acquisition across acquisition dates, we adopt the approach developed by Wooldridge (2023).[5] This method consists in grouping observations in 'cohorts' for which the treatment occurs at the same time, and then estimating the treatment effects for each cohort separately.

We first define 'acquisition-cohorts', i.e. patents acquired at the same time. For the parameter estimates of the model to converge, we need enough variability (and hence a sufficient number of patents) in each cohort. On this basis, the smallest cohort size that can be used is 7 months, with a study period starting in February 2006. So we group all patents acquired within the same 7 months. We obtain 16 cohorts (C = 16), and each

[5]The estimation method and the underlying assumptions are presented on pages C46 to C49 of Wooldridge (2023).

cohort contains on average 41 patents.

We then estimate the treatment effects for each cohort separately:

$$Cit_{i,c} = f(I'\theta^7, Cohort'\zeta^4, I'Cohort'\eta^4, N'\gamma^7) + \varepsilon_{i,c} \tag{A.6}$$

where I' is a vector containing the time dummies at the 7 months-level, $Cohort'$ contains cohort dummies for patents acquired within the same 7 months, and N' contains the calendar date dummies grouped over 7 months.

Finally, we take the weighted average of the cohort-specific treatment effects, where the weights (w_c) indicate the number of patents in each cohort c. This defines the average treatment effect in each event time i:

$$ATT_i = \frac{1}{C} \sum_c w_c \; (e^{\theta_i^7 + \eta_{i,c}^4}) \tag{A.7}$$

Figure A.9: Big Tech citations to acquired patents relative to acquisition, by cohort

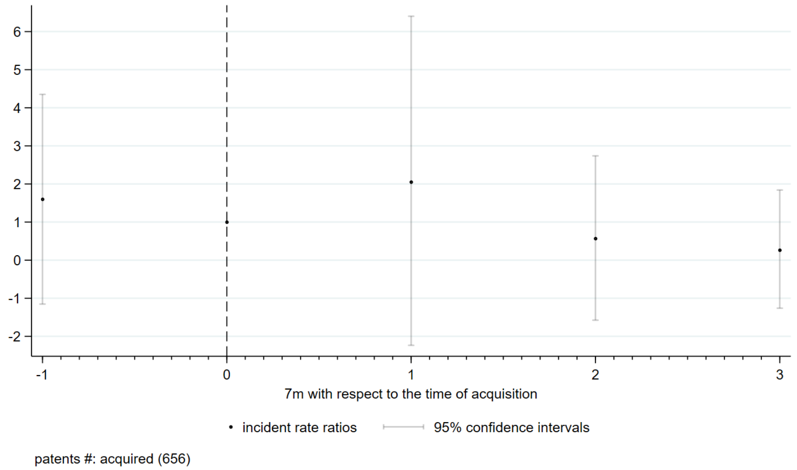

Notes: The graph shows the incident rate ratios for acquired patents aggregated across all acquisition-cohorts: ATT_i from equation A.7. These coefficients are estimated on a balanced sample of patents in a 3 year-window around acquisition.

From Figure A.9, and despite much larger confidence intervals due to smaller (cohort) sample sizes, we can see that our findings are robust to allowing for heterogeneous treatment effects; we observe a non-lasting boost in citations after acquisition.

Appendix B

B.1 Distribution of Orbis firm-level data

Figure B.1: Number of employees at acquisition (boxplot)

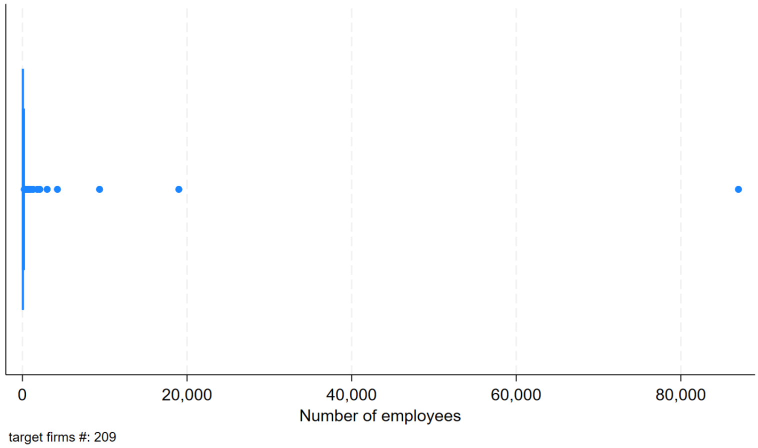

Note: The three outlier observations are LINKEDIN (9, 372 employees), MOTOROLA MOBILITY (19, 000 employees) and WHOLE FOODS MARKET (87, 000 employees).

Figure B.2: Funding amount at acquisition (boxplot)

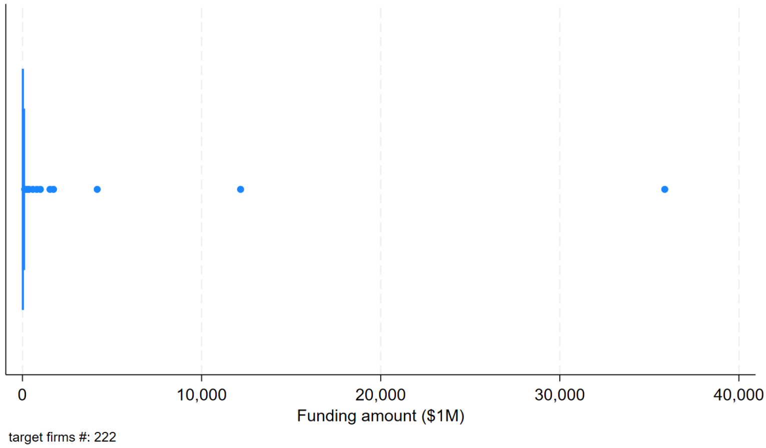

Note: The three outlier observations are FITBIT ($4.17 billion), LINKEDIN ($12.18 billion) and WHOLE FOODS MARKET ($35.86 billion).

B.2 Descriptive statistics on USPTO patent-level data

Table B.1: Statistics over all the patents filed by Big Tech-acquired inventors

	US patents		Core patents		Patents market value		Co-authored patents		Filing dates	
	obs.	count (%)	obs.	count (%)	obs.	mean (SD)	obs.	count (%)	obs.	[min,max]
AMZN	588	498 (85%)	527	341 (65%)	460	-.01 (.03)	596	499 (84%)	596	1995m6, 2020m6
APPL	700	510 (73%)	724	280 (39%)	714	-.01 (.03)	730	613 (84%)	739	1988m6, 2019m11
FCBK	55	44 (80%)	133	108 (81%)	132	-.02 (.03)	136	114 (84%)	136	1998m11, 2019m1
GOOG	2057	1859 (90%)	1950	465 (24%)	1932	-.02 (.04)	2076	1783 (86%)	2082	1995m6, 2020m9
MSFT	1439	1027 (71%)	1430	1151 (80%)	1407	-.01 (.04)	1452	1022 (70%)	1454	1987m5, 2020m4
TOTAL	4839	3938 (81%)	4764	2345 (49%)	4645	-.02 (.04)	4990	4031 (81%)	5007	1987m5, 2020m9

This table presents statistics at the patent-level.

B.3 Core variable aggregated at the target level

We can assess whether a target T is active in its acquirer's core technology fields:

$$Core_T = max\ I_{t,T}$$

where $I_{t,T} = 1$ if a given patent t of the target's portfolio is associated with at least one of its acquirer's core fields in the year this patent was acquired.

Table B.2: Big Tech acquired patents portfolios

	Count	$\text{Core}_T = 1$
AMZN	27	19 (70%)
APPL	52	27 (52%)
FCBK	17	10 (59%)
GOOG	67	40 (60%)
MSFT	88	81 (92%)
TOTAL	251	177 (71%)

From Table B.2, we see that 71% of Big Tech targets own at least one patent belonging to their acquirer's core business.

B.4 Components of the market value index

We construct a market value index using indicators of a patent's economic value: family size, grant lag and number of claims. For each of these variables, we explain here, based on the OECD note accompanying the related data (Squicciarini, Dernis, and Criscuolo 2013), why they are good measures of a patent's economic value.

Family size

A patent family is a set of patents filed in several countries but with a common priority filing. The size of patent families is proxied by the number of patent offices at which a given invention has been protected. Because extending a patent protection to other countries implies additional costs and delays, applicants are more likely to go through that procedure for more valuable patents.

Grant lag

The grant lag variable is constructed based on the number of days elapsing between the patent application and granting date. In line with the argument that applicants try to accelerate the grant procedure for their most valuable patents, the length of the grant lag period has been shown to be negatively correlated with the value of a patent (Harhoff and Wagner 2009; Régibeau and Rockett 2010).

Adjusted number of claims

A patent is composed of claims, which relate to the technologies that are legally protected by the patent. Therefore, the more claims a patent contains, the broader the rights conferred by this patent. It has been shown that patents containing more claims have, on average, a higher market value (Tong and Frame 1994; Lanjouw and Schankerman 2001, 2004). Because technology fields seem to vary in the average number of claims per patent, this variable is further adjusted. The number of citations to prior art by a patent is used to account for the development level of the technology area to which this patent belongs, and the adjusted variable is defined as the number of claims over the number of citations.

B.5 Market Value Index

As a measure of innovation market value, we propose to define one vector that approximates the information contained in three (or, in an alternative specification, four) distinct indicators: family size, grant lag, adjusted number of claims (and forward citations). To do so, we follow a Principal Component Analysis. First, all these variables are normalized such as to be centered in 0:

$$\widetilde{X} = [(X - min(X))\,/(max(X) - min(X))] - mean(X)$$

Next, we center our observations around a point that has as coordinates the average values of all three (normalized) indicators. The vector that maximises the variation in the data along the three market value indicators is calculated as the line that best fits the data while going through this center.

Table B.3: Market value indicators

	Weights (ev^x)	
Family size	0.49	0.52
Grant lag	-0.64	-0.49
Adjusted claims #	0.59	0.43
Forward citations #		0.55
Eigenvalue Proportion	37%	29%
Patents #	7,031,531	5,962,352

Note: This table presents the eigenvectors from the Principal Component Analysis eigen decomposition.

The obtained vector is defined by three values, each of which capturing the importance of the associated market value indicator in positioning the vector in space (see Table B.3). We obtain the values of the principal components in our sample by projecting our data onto this vector. Because changing the signs of the components does not change the variance that they contain (we are just projecting on a vector that is pointing in one direction or 180° away in the other direction), we can define the signs of the indicators based on the relation the associated metric is expected to have with market value: positive for family size, number of claims and forward citations, negative for grant lag. Patents associated with higher values of this projection can thus be considered as more valuable.

The proportion of the variation in the data that the market value index accounts for can be computed as the distance between the data projections and the center of the data. This distance is called the eigenvalue of a component. As we can see from Table B.3, the projection of the three market value indicators account for 37% of the variation in the data. Including forward citations does not seem to improve the performance of the index, since only 29% of the variation in the data would then be captured.

B.6 Descriptive statistics on the Talent variable

Table B.4: Big Tech acquired patents portfolios

	Count	$\sum_i Talent_i > 1$	$\sum_i Talent^*{}_i > 1$
AMZN	27	21 (78%)	15 (56%)
APPL	52	40 (77%)	30 (58%)
FCBK	18	11 (61%)	8 (44%)
GOOG	67	50 (75%)	28 (42%)
MSFT	88	64 (73%)	43 (49%)
TOTAL	252	186 (74%)	124 (49%)

On Figure B.3, we compare the talent indices normalized by the total number of inventors working at the target T and at the acquirer A:

$$Talent_{A,T} = \frac{log(\sum_i \{Talent_i\} + e^1)}{InventorsNbr_T + InventorsNbr_A}$$

where the numerator is log-transformed such as to reduce the influence of outliers.

Figure B.3: Normalized Talent indices, by acquirer – Distribution

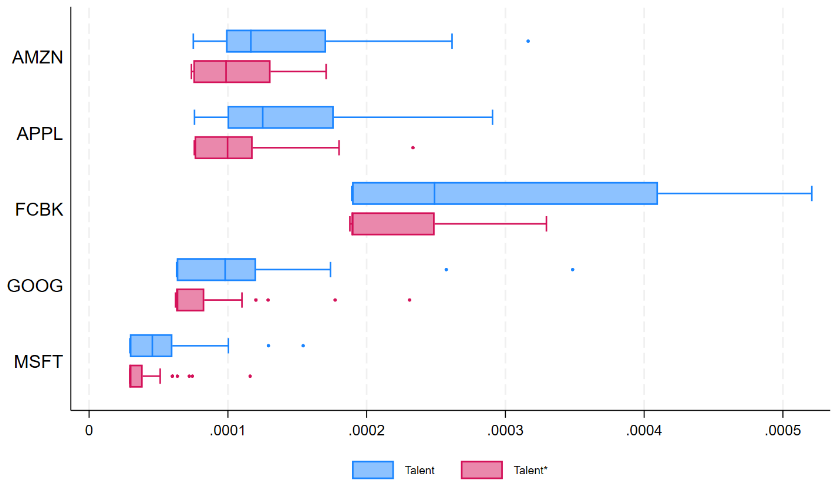

B.7 Model 2.1 – Estimation results for *Talent**

Table B.5: Inventors innovating for their acquirer with some acquirer's employees

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent*									
Core	0.371***	0.335**	0.379***	0.344**	0.340**	0.433**	0.392*	0.536**	0.461*
	(2.62)	(2.35)	(2.59)	(2.34)	(2.31)	(2.12)	(1.91)	(2.13)	(1.74)
MSinceAcqui	0.009***	0.016***	0.016***	0.008***	0.009***	0.016***	0.016***	0.022***	0.024***
	(10.13)	(8.31)	(7.68)	(3.07)	(3.54)	(5.03)	(4.98)	(5.79)	(5.65)
PatentsCount	0.030***	0.037***	0.034***	0.020***	0.018**	0.016*	0.016*	0.016*	0.017*
	(4.38)	(4.88)	(4.66)	(2.93)	(2.50)	(1.85)	(1.91)	(1.67)	(1.83)
FirstFiling		0.080***	0.078***	-0.026	-0.008	0.014	0.013	0.044	0.046
		(4.02)	(3.66)	(-0.94)	(-0.27)	(0.42)	(0.39)	(1.15)	(1.17)
US			0.255*	0.281*	0.285*	0.250	1.387***	1.277**	-0.396
			(1.70)	(1.86)	(1.88)	(1.28)	(2.67)	(2.27)	(-0.37)
MSinceLastFil				-0.015***	-0.013***	-0.013***	-0.013***	-0.015***	-0.016***
				(-5.09)	(-4.48)	(-3.75)	(-3.76)	(-3.80)	(-3.78)
FirstAuthor					0.411***	0.393***	0.396***	0.373**	0.360**
					(3.56)	(2.97)	(2.99)	(2.52)	(2.36)
Incorp^T						0.017	0.015	-0.011	0.003
						(0.98)	(0.86)	(-0.57)	(0.15)
US^T							-1.157**	-1.112**	0.126
							(-2.35)	(-2.13)	(0.12)
FirmSize^T								-0.509***	-0.761***
								(-4.26)	(-4.42)
Funding^T									0.427**
									(1.98)
Constant	-3.530***	-166.007***	-160.365***	48.670	11.906	-65.652	-59.635	-71.341	-102.684
	(-17.53)	(-4.11)	(-3.75)	(0.88)	(0.21)	(-0.94)	(-0.85)	(-0.90)	(-1.22)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

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

B.8 Model 2.1 – Robustness checks

Table B.6: Innovating for acquirer, fixed period

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core	0.455***	0.413***	0.331***	0.343***	0.345***	0.642***	0.664***	0.779***	0.958***
	(4.47)	(3.87)	(3.00)	(3.09)	(3.10)	(4.36)	(4.46)	(3.91)	(4.46)
MSinceAcqui	0.001	0.021***	0.022***	0.008***	0.008***	0.009***	0.009***	0.007**	0.007*
	(1.07)	(12.71)	(12.73)	(3.40)	(3.44)	(3.36)	(3.33)	(2.11)	(1.84)
PatentsCount	0.251***	0.339***	0.340***	0.214***	0.198***	0.156***	0.154***	0.139***	0.141***
	(13.72)	(16.24)	(16.04)	(9.15)	(8.49)	(6.82)	(6.79)	(6.04)	(6.00)
FirstFiling		0.218***	0.222***	0.053**	0.057**	0.030	0.029	−0.002	−0.002
		(14.18)	(13.97)	(2.14)	(2.29)	(1.15)	(1.13)	(−0.07)	(−0.08)
US			−0.325***	−0.317***	−0.312***	−0.317**	−0.783*	−0.091	−0.221
			(−3.12)	(−3.00)	(−2.94)	(−2.23)	(−1.73)	(−0.17)	(−0.34)
MSinceLastFil				−0.021***	−0.021***	−0.024***	−0.024***	−0.030***	−0.031***
				(−8.37)	(−8.32)	(−8.72)	(−8.76)	(−9.66)	(−9.93)
FirstAuthor					0.293***	0.331***	0.332***	0.268***	0.252***
					(3.81)	(3.91)	(3.92)	(2.91)	(2.68)
Incorp^{T}						0.035**	0.037**	0.009	−0.003
						(2.39)	(2.52)	(0.44)	(−0.12)
US^{T}							0.476	−0.310	−0.183
							(1.09)	(−0.58)	(−0.28)
FirmSize^{T}								−0.428***	−0.353**
								(−4.33)	(−2.41)
Funding^{T}									−0.029
									(−0.17)
Constant	−0.578***	−440.761***	−449.468***	−107.060**	−114.695**	−129.441**	−133.155**	−13.223	10.247
	(−3.56)	(−14.20)	(−13.98)	(−2.14)	(−2.29)	(−2.23)	(−2.29)	(−0.19)	(0.14)
Observations	3946	3946	3852	3852	3852	3240	3240	2861	2787
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.7: Innovating for acquirer with some acquirer's employees, fixed period

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent*									
Core	0.342**	0.301*	0.362**	0.345**	0.338**	0.426*	0.420*	0.590**	0.512*
	(2.13)	(1.87)	(2.19)	(2.09)	(2.04)	(1.88)	(1.83)	(2.05)	(1.66)
MSinceAcqui	0.004***	0.010***	0.011***	0.003	0.004	0.008**	0.008**	0.015***	0.023***
	(3.53)	(4.48)	(4.56)	(1.05)	(1.31)	(2.04)	(2.04)	(3.35)	(4.32)
PatentsCount	0.072***	0.079***	0.075***	0.051***	0.045***	0.041***	0.041***	0.036**	0.040***
	(6.94)	(7.33)	(6.86)	(4.41)	(3.87)	(2.88)	(2.88)	(2.37)	(2.60)
FirstFiling		0.067***	0.074***	−0.021	−0.009	0.015	0.015	0.064	0.072
		(3.07)	(3.29)	(−0.72)	(−0.30)	(0.41)	(0.41)	(1.50)	(1.62)
US			0.503***	0.545***	0.556***	0.435*	0.555	0.400	0.004
			(2.94)	(3.16)	(3.22)	(1.94)	(0.91)	(0.59)	(0.00)
MSinceLastFil				−0.014***	−0.013***	−0.012***	−0.012***	−0.013***	−0.013***
				(−4.27)	(−3.91)	(−3.24)	(−3.24)	(−2.86)	(−2.78)
FirstAuthor					0.357***	0.383***	0.383***	0.389**	0.383**
					(2.85)	(2.64)	(2.64)	(2.47)	(2.36)
Incorp^{T}						0.036	0.035	0.021	0.050
						(1.62)	(1.55)	(0.71)	(1.47)
US^{T}							−0.124	−0.159	−0.567
							(−0.21)	(−0.24)	(−0.51)
FirmSize^{T}								−0.337**	−0.960***
								(−2.23)	(−4.13)
Funding^{T}									1.052***
									(3.72)
Constant	−2.614***	−136.997***	−153.097***	40.563	15.686	−104.822	−102.801	−172.135*	−248.347**
	(−11.12)	(−3.13)	(−3.35)	(0.67)	(0.26)	(−1.30)	(−1.27)	(−1.77)	(−2.36)
Observations	3946	3946	3852	3852	3852	3240	3240	2861	2787
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

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

Table B.8: Innovating for acquirer, with buffer

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core	0.411***	0.364***	0.324***	0.297***	0.297***	0.525***	0.512***	0.718***	0.770***
	(4.80)	(4.05)	(3.54)	(3.18)	(3.16)	(4.25)	(4.12)	(4.53)	(4.52)
MSinceAcqui	0.013***	0.031***	0.030***	0.014***	0.014***	0.018***	0.018***	0.023***	0.024***
	(16.75)	(20.36)	(19.45)	(7.38)	(7.81)	(8.44)	(8.43)	(8.85)	(9.05)
PatentsCount	0.055***	0.089***	0.089***	0.028***	0.024***	0.024***	0.024***	0.022***	0.021***
	(7.71)	(10.58)	(10.49)	(3.65)	(3.18)	(2.89)	(2.91)	(2.72)	(2.64)
FirstFiling		0.197***	0.191***	-0.021	-0.010	-0.025	-0.025	-0.050**	-0.057**
		(14.62)	(13.61)	(-1.12)	(-0.51)	(-1.16)	(-1.16)	(-2.11)	(-2.31)
US			-0.280***	-0.245***	-0.246***	-0.246*	0.151	0.268	-0.263
			(-3.21)	(-2.73)	(-2.73)	(-1.89)	(0.38)	(0.59)	(-0.46)
MSinceLastFil				-0.031***	-0.030***	-0.031***	-0.031***	-0.035***	-0.037***
				(-14.62)	(-14.17)	(-13.28)	(-13.26)	(-13.32)	(-13.47)
FirstAuthor					0.292***	0.331***	0.332***	0.274***	0.279***
					(4.13)	(4.21)	(4.22)	(3.20)	(3.19)
Incorp^T						0.021**	0.021**	-0.001	-0.001
						(2.15)	(2.12)	(-0.07)	(-0.08)
US^T							-0.404	-0.272	0.016
							(-1.06)	(-0.64)	(0.03)
FirmSize^T								-0.541***	-0.689***
								(-6.76)	(-5.74)
Funding^T									0.237*
									(1.78)
Constant	-1.907***	-399.599***	-387.719***	42.031	18.405	5.990	6.610	101.278**	114.766**
	(-15.40)	(-14.69)	(-13.67)	(1.09)	(0.47)	(0.13)	(0.14)	(1.99)	(2.17)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.9: Innovating for acquirer with some acquirer's employees, with buffer

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent*									
Core	0.366**	0.329**	0.368**	0.331**	0.328**	0.450**	0.406*	0.614**	0.553**
	(2.56)	(2.28)	(2.49)	(2.23)	(2.20)	(2.17)	(1.94)	(2.40)	(2.04)
MSinceAcqui	0.010***	0.017***	0.016***	0.008***	0.009***	0.017***	0.017***	0.022***	0.024***
	(10.28)	(8.33)	(7.76)	(3.07)	(3.54)	(5.09)	(5.04)	(5.89)	(5.75)
PatentsCount	0.031***	0.037***	0.035***	0.021***	0.018**	0.016*	0.017*	0.016	0.017*
	(4.41)	(4.90)	(4.70)	(2.94)	(2.52)	(1.88)	(1.95)	(1.63)	(1.80)
FirstFiling		0.080***	0.078***	-0.028	-0.010	0.013	0.012	0.043	0.045
		(3.97)	(3.67)	(-1.02)	(-0.35)	(0.40)	(0.37)	(1.13)	(1.16)
US			0.214	0.241	0.246	0.235	1.508***	1.347**	-0.274
			(1.42)	(1.59)	(1.62)	(1.19)	(2.88)	(2.37)	(-0.25)
MSinceLastFil				-0.016***	-0.014***	-0.013***	-0.013***	-0.016***	-0.016***
				(-5.23)	(-4.63)	(-3.92)	(-3.92)	(-3.90)	(-3.90)
FirstAuthor					0.411***	0.405***	0.408***	0.389***	0.379**
					(3.53)	(3.02)	(3.04)	(2.60)	(2.46)
Incorp^T						0.013	0.011	-0.013	0.003
						(0.78)	(0.64)	(-0.68)	(0.12)
US^T							-1.296***	-1.185**	-0.020
							(-2.61)	(-2.24)	(-0.02)
FirmSize^T								-0.493***	-0.757***
								(-4.06)	(-4.34)
Funding^T									0.454**
									(2.06)
Constant	-3.622***	-165.770***	-162.083***	53.131	16.497	-57.509	-50.621	-65.914	-101.421
	(-17.56)	(-4.06)	(-3.76)	(0.96)	(0.29)	(-0.82)	(-0.72)	(-0.83)	(-1.19)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

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

Table B.10: Innovating under acquirer's name

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core	0.493***	0.458***	0.473***	0.441***	0.441***	0.855***	0.828***	1.003***	1.076***
	(5.56)	(4.99)	(5.03)	(4.60)	(4.59)	(6.64)	(6.40)	(5.91)	(5.77)
MSinceAcqui	0.011***	0.025***	0.025***	0.010***	0.010***	0.014***	0.014***	0.017***	0.019***
	(14.92)	(17.28)	(16.66)	(5.52)	(5.87)	(6.55)	(6.53)	(6.65)	(7.14)
PatentsCount	0.046***	0.069***	0.067***	0.018***	0.015**	0.011	0.011*	0.010	0.009
	(6.83)	(9.09)	(8.71)	(2.66)	(2.30)	(1.61)	(1.67)	(1.43)	(1.28)
FirstFiling		0.154***	0.157***	−0.038**	−0.028	−0.037*	−0.037*	−0.056**	−0.064***
		(11.67)	(11.14)	(−2.01)	(−1.48)	(−1.73)	(−1.74)	(−2.38)	(−2.60)
US			−0.089	−0.051	−0.050	0.161	1.068***	1.905***	0.891
			(−0.99)	(−0.56)	(−0.54)	(1.21)	(2.63)	(3.88)	(1.45)
MSinceLastFil				−0.029***	−0.028***	−0.028***	−0.028***	−0.033***	−0.035***
				(−13.78)	(−13.34)	(−12.29)	(−12.27)	(−12.56)	(−12.77)
FirstAuthor					0.230***	0.302***	0.304***	0.245***	0.266***
					(3.18)	(3.74)	(3.75)	(2.74)	(2.91)
$Incorp^T$						0.018*	0.017	−0.005	0.009
						(1.75)	(1.64)	(−0.38)	(0.59)
US^T							−0.920**	−1.478***	−0.999*
							(−2.36)	(−3.17)	(−1.70)
$FirmSize^T$								−0.556***	−0.880***
								(−6.71)	(−6.93)
$Funding^T$									0.524***
									(3.62)
Constant	−2.370***	−313.216***	−318.900***	74.151*	54.792	33.702	36.144	118.828**	106.533*
	(−17.86)	(−11.76)	(−11.22)	(1.96)	(1.43)	(0.74)	(0.79)	(2.33)	(1.95)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.11: Innovating under acquirer's name with some acquirer's employees

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent*									
Core	0.404***	0.365**	0.411***	0.377**	0.374**	0.463**	0.421**	0.553**	0.476*
	(2.82)	(2.53)	(2.77)	(2.54)	(2.50)	(2.24)	(2.02)	(2.18)	(1.78)
MSinceAcqui	0.009***	0.017***	0.016***	0.008***	0.010***	0.017***	0.017***	0.022***	0.024***
	(9.90)	(8.34)	(7.72)	(3.25)	(3.75)	(5.07)	(5.02)	(5.74)	(5.55)
PatentsCount	0.031***	0.038***	0.035***	0.022***	0.019***	0.017**	0.018**	0.017*	0.018**
	(4.39)	(4.92)	(4.70)	(3.06)	(2.63)	(2.02)	(2.08)	(1.83)	(1.97)
FirstFiling		0.086***	0.084***	−0.017	0.003	0.023	0.022	0.052	0.054
		(4.21)	(3.87)	(−0.60)	(0.10)	(0.67)	(0.64)	(1.33)	(1.36)
US			0.262*	0.287*	0.292*	0.265	1.458***	1.365**	−0.331
			(1.72)	(1.87)	(1.90)	(1.33)	(2.80)	(2.41)	(−0.31)
MSinceLastFil				−0.015***	−0.013***	−0.012***	−0.012***	−0.015***	−0.016***
				(−4.85)	(−4.21)	(−3.52)	(−3.53)	(−3.70)	(−3.71)
FirstAuthor					0.444***	0.445***	0.448***	0.443***	0.429***
					(3.79)	(3.30)	(3.32)	(2.93)	(2.75)
$Incorp^T$						0.022	0.019	−0.008	0.004
						(1.24)	(1.12)	(−0.43)	(0.18)
US^T							−1.214**	−1.201**	0.096
							(−2.46)	(−2.29)	(0.09)
$FirmSize^T$								−0.532***	−0.761***
								(−4.39)	(−4.35)
$Funding^T$									0.389*
									(1.79)
Constant	−3.618***	−177.646***	−173.335***	30.316	−9.735	−93.429	−86.997	−92.224	−122.266
	(−17.50)	(−4.30)	(−3.96)	(0.54)	(−0.17)	(−1.30)	(−1.21)	(−1.13)	(−1.42)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

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

Table B.12: Innovating for acquirer, core w/i 1y

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core	0.897***	0.526***	0.485***	0.123	0.129	0.093	0.093	0.219*	0.212
	(10.34)	(5.72)	(5.20)	(1.26)	(1.31)	(0.82)	(0.82)	(1.72)	(1.62)
MSinceAcqui	0.012***	0.027***	0.027***	0.012***	0.012***	0.016***	0.016***	0.018***	0.019***
	(16.29)	(18.51)	(17.76)	(6.44)	(6.84)	(7.46)	(7.47)	(7.58)	(7.64)
PatentsCount	0.065***	0.100***	0.099***	0.042***	0.037***	0.039***	0.039***	0.035***	0.035***
	(8.41)	(11.16)	(10.94)	(4.89)	(4.30)	(4.05)	(4.07)	(3.66)	(3.62)
FirstFiling		0.168***	0.168***	-0.021	-0.010	-0.020	-0.020	-0.031	-0.037
		(12.58)	(11.90)	(-1.10)	(-0.52)	(-0.91)	(-0.91)	(-1.30)	(-1.52)
US			-0.296***	-0.256***	-0.256***	-0.346***	0.051	-0.106	-0.806
			(-3.43)	(-2.91)	(-2.89)	(-2.76)	(0.13)	(-0.24)	(-1.44)
MSinceLastFil				-0.029***	-0.028***	-0.029***	-0.029***	-0.033***	-0.035***
				(-13.32)	(-12.88)	(-12.12)	(-12.10)	(-12.28)	(-12.38)
FirstAuthor					0.295***	0.328***	0.329***	0.270***	0.267***
					(4.23)	(4.25)	(4.26)	(3.21)	(3.11)
Incorp^T						0.018**	0.018**	-0.007	-0.009
						(2.08)	(2.03)	(-0.75)	(-0.87)
US^T							-0.401	-0.175	0.321
							(-1.06)	(-0.42)	(0.60)
FirmSize^T								-0.629***	-0.731***
								(-8.29)	(-6.63)
Funding^T									0.170
									(1.42)
Constant	-1.783***	-341.909***	-341.571***	41.696	19.245	2.407	2.990	76.934	92.799*
	(-16.15)	(-12.64)	(-11.95)	(1.07)	(0.49)	(0.05)	(0.07)	(1.55)	(1.82)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.13: Innovating for acquirer with some acquirer's employees, core w/i 1y

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent*									
Core	0.712***	0.576***	0.539***	0.333**	0.331**	0.417**	0.422**	0.633***	0.668***
	(5.23)	(4.01)	(3.69)	(2.19)	(2.17)	(2.28)	(2.31)	(3.03)	(3.11)
MSinceAcqui	0.010***	0.015***	0.014***	0.007***	0.009***	0.016***	0.016***	0.021***	0.023***
	(10.40)	(7.39)	(6.90)	(3.00)	(3.47)	(4.93)	(4.92)	(5.66)	(5.56)
PatentsCount	0.028***	0.033***	0.032***	0.022***	0.019***	0.017*	0.018**	0.017	0.018*
	(4.37)	(4.73)	(4.59)	(3.11)	(2.65)	(1.94)	(2.00)	(1.63)	(1.79)
FirstFiling		0.058***	0.059***	-0.025	-0.006	0.014	0.014	0.045	0.046
		(2.86)	(2.73)	(-0.90)	(-0.23)	(0.43)	(0.41)	(1.18)	(1.17)
US			0.238	0.263*	0.269*	0.212	1.483***	1.357**	-0.455
			(1.59)	(1.74)	(1.78)	(1.10)	(2.86)	(2.40)	(-0.42)
MSinceLastFil				-0.013***	-0.012***	-0.010***	-0.010***	-0.012***	-0.012***
				(-4.28)	(-3.72)	(-2.88)	(-2.87)	(-2.98)	(-2.87)
FirstAuthor					0.412***	0.393***	0.397***	0.375**	0.362**
					(3.57)	(2.97)	(2.99)	(2.53)	(2.37)
Incorp^T						0.023	0.019	-0.009	0.005
						(1.36)	(1.16)	(-0.51)	(0.22)
US^T							-1.283***	-1.270**	0.120
							(-2.63)	(-2.43)	(0.11)
FirmSize^T								-0.589***	-0.822***
								(-5.31)	(-5.02)
Funding^T									0.430**
									(1.98)
Constant	-3.552***	-121.456***	-122.038***	46.306	9.548	-77.997	-70.419	-76.396	-106.287
	(-18.96)	(-2.94)	(-2.81)	(0.84)	(0.17)	(-1.11)	(-1.00)	(-0.96)	(-1.26)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

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

Table B.14: Innovating for acquirer, 3-digit core

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core	0.392***	0.360***	0.320***	0.276***	0.272***	0.480***	0.474***	0.507***	0.484***
	(4.43)	(3.95)	(3.43)	(2.90)	(2.85)	(4.17)	(4.10)	(3.66)	(3.37)
MSinceAcqui	0.012***	0.029***	0.029***	0.012***	0.013***	0.017***	0.017***	0.020***	0.021***
	(15.78)	(20.00)	(19.12)	(6.80)	(7.18)	(8.03)	(8.03)	(8.12)	(8.14)
PatentsCount	0.070***	0.107***	0.105***	0.040***	0.035***	0.036***	0.036***	0.033***	0.033***
	(9.16)	(11.91)	(11.62)	(4.72)	(4.16)	(3.79)	(3.81)	(3.50)	(3.47)
FirstFiling		0.190***	0.188***	-0.020	-0.009	-0.018	-0.018	-0.035	-0.040
		(14.60)	(13.69)	(-1.06)	(-0.49)	(-0.84)	(-0.83)	(-1.44)	(-1.61)
US			-0.266***	-0.234***	-0.234***	-0.283**	-0.019	-0.111	-0.683
			(-3.11)	(-2.65)	(-2.65)	(-2.25)	(-0.05)	(-0.25)	(-1.22)
MSinceLastFil				-0.030***	-0.029***	-0.029***	-0.029***	-0.034***	-0.036***
				(-14.14)	(-13.70)	(-12.71)	(-12.70)	(-12.89)	(-13.06)
FirstAuthor					0.292***	0.332***	0.332***	0.269***	0.265***
					(4.18)	(4.28)	(4.29)	(3.20)	(3.08)
Incorp^T						0.014*	0.014	-0.008	-0.010
						(1.65)	(1.62)	(-0.76)	(-0.94)
US^T							-0.268	-0.070	0.336
							(-0.71)	(-0.17)	(0.63)
FirmSize^T								-0.560***	-0.646***
								(-7.21)	(-5.70)
Funding^T									0.137
									(1.14)
Constant	-1.805***	-385.519***	-381.985***	40.299	17.974	5.870	6.196	83.754*	98.801*
	(-14.23)	(-14.67)	(-13.74)	(1.04)	(0.46)	(0.13)	(0.14)	(1.68)	(1.93)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table B.15: Innovating for acquirer with some acquirer's employees, 3-digit core

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent⋆									
Core	0.475***	0.447***	0.472***	0.420***	0.411***	0.519***	0.490**	0.350	0.294
	(3.26)	(3.06)	(3.17)	(2.82)	(2.74)	(2.74)	(2.56)	(1.55)	(1.24)
MSinceAcqui	0.010***	0.017***	0.016***	0.008***	0.009***	0.017***	0.017***	0.022***	0.023***
	(10.35)	(8.49)	(7.88)	(3.26)	(3.72)	(5.25)	(5.20)	(5.75)	(5.62)
PatentsCount	0.030***	0.037***	0.034***	0.021***	0.018**	0.016*	0.017*	0.017*	0.018*
	(4.36)	(4.86)	(4.67)	(2.94)	(2.53)	(1.90)	(1.95)	(1.80)	(1.94)
FirstFiling		0.081***	0.079***	-0.023	-0.005	0.017	0.016	0.046	0.047
		(4.04)	(3.72)	(-0.84)	(-0.17)	(0.52)	(0.49)	(1.21)	(1.21)
US			0.244	0.269*	0.273*	0.224	1.362***	1.228**	-0.474
			(1.63)	(1.79)	(1.81)	(1.16)	(2.62)	(2.19)	(-0.44)
MSinceLastFil				-0.015***	-0.013***	-0.012***	-0.012***	-0.015***	-0.016***
				(-5.00)	(-4.39)	(-3.62)	(-3.63)	(-3.80)	(-3.79)
FirstAuthor					0.407***	0.389***	0.393***	0.369**	0.357**
					(3.52)	(2.94)	(2.96)	(2.49)	(2.34)
Incorp^T						0.019	0.017	-0.006	0.007
						(1.17)	(1.03)	(-0.32)	(0.31)
US^T							-1.153**	-1.132**	0.122
							(-2.35)	(-2.17)	(0.11)
FirmSize^T								-0.543***	-0.801***
								(-4.64)	(-4.73)
Funding^T									0.441**
									(2.09)
Constant	-3.639***	-166.450***	-162.413***	42.629	5.852	-77.458	-70.980	-85.315	-112.618
	(-17.39)	(-4.13)	(-3.81)	(0.77)	(0.10)	(-1.11)	(-1.02)	(-1.08)	(-1.35)
Observations	4966	4966	4806	4806	4806	4051	4051	3593	3519
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

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

B.9 Observational cut in the patents database

Figure B.4: Talent dummy over acquisition year

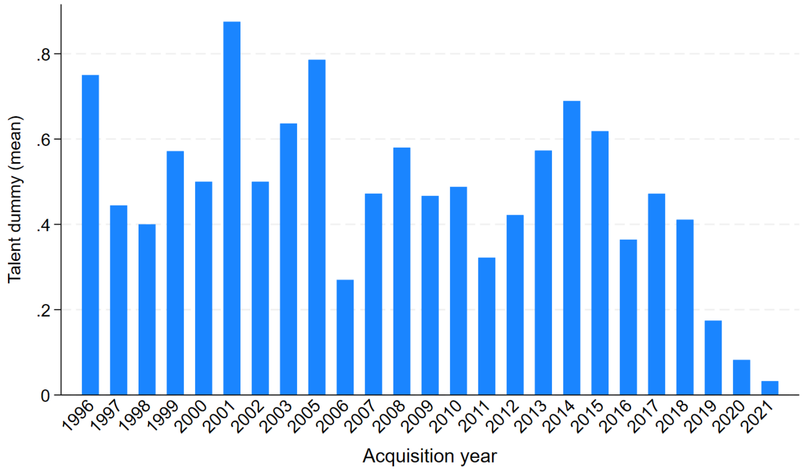

B.10 Model 2.2 – Alternative Market Value Indices

Table B.16: Inventors innovating for their acquirer, MarketValue definition (a)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core=1	0.044	−0.077	−0.126	−0.421*	−0.406	−0.340	−0.343	−0.174	−0.180
	(0.19)	(−0.33)	(−0.52)	(−1.66)	(−1.60)	(−1.17)	(−1.18)	(−0.52)	(−0.53)
MarketVal_bar	0.046	−0.820	−1.032*	−1.905***	−1.848***	−1.107	−1.109*	−2.023***	−2.208***
	(0.09)	(−1.54)	(−1.92)	(−3.28)	(−3.17)	(−1.64)	(−1.65)	(−2.77)	(−2.94)
Core=1 × MarketVal_bar	1.192*	1.446*	1.478*	2.414***	2.359***	2.669***	2.661***	2.319**	2.577**
	(1.65)	(1.94)	(1.95)	(3.02)	(2.94)	(2.91)	(2.90)	(2.26)	(2.46)
MSinceAcqui	0.012***	0.029***	0.029***	0.012***	0.013***	0.016***	0.016***	0.019***	0.020***
	(15.50)	(19.80)	(18.94)	(6.74)	(7.11)	(7.69)	(7.68)	(7.82)	(7.89)
PatentsCount	0.068***	0.105***	0.104***	0.038***	0.033***	0.033***	0.033***	0.032***	0.032***
	(8.82)	(11.69)	(11.41)	(4.53)	(3.99)	(3.56)	(3.57)	(3.37)	(3.35)
FirstFiling		0.193***	0.193***	−0.019	−0.008	−0.019	−0.019	−0.030	−0.034
		(14.62)	(13.72)	(−0.98)	(−0.44)	(−0.87)	(−0.87)	(−1.25)	(−1.35)
US			−0.239***	−0.212**	−0.212**	−0.258**	−0.098	−0.179	−0.661
			(−2.74)	(−2.36)	(−2.36)	(−2.01)	(−0.25)	(−0.40)	(−1.17)
MSinceLastFil				−0.030***	−0.030***	−0.031***	−0.031***	−0.036***	−0.037***
				(−14.47)	(−14.04)	(−13.04)	(−13.03)	(−13.25)	(−13.41)
FirstAuthor					0.278***	0.320***	0.320***	0.256***	0.248***
					(3.96)	(4.10)	(4.11)	(3.02)	(2.87)
$Incorp^T$						0.012	0.012	−0.011	−0.017
						(1.25)	(1.24)	(−1.08)	(−1.52)
US^T							−0.163	−0.006	0.442
							(−0.43)	(−0.01)	(0.83)
$FirmSize^T$								−0.583***	−0.598***
								(−7.36)	(−5.07)
$Funding^T$									0.044
									(0.35)
Constant	−1.717***	−392.295***	−390.551***	37.879	16.627	13.516	13.611	83.648*	102.000**
	(−8.86)	(−14.68)	(−13.77)	(0.97)	(0.42)	(0.29)	(0.30)	(1.65)	(1.96)
Observations	4880	4880	4720	4720	4720	3967	3967	3509	3435
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

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

Table B.17: Inventors innovating for their acquirer, MarketValue definition (b.2)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core=1	0.456***	0.433***	0.403***	0.448***	0.442***	0.653***	0.648***	0.655***	0.748***
	(4.87)	(4.44)	(4.08)	(4.42)	(4.36)	(4.97)	(4.91)	(4.18)	(4.46)
MarketVal_bis	3.466**	-1.137	-1.878	-6.184***	-5.854***	-5.638***	-5.647***	-9.002***	-9.602***
	(2.48)	(-0.76)	(-1.23)	(-3.74)	(-3.53)	(-3.12)	(-3.12)	(-4.47)	(-4.64)
Core=1 × MarketVal_bis	4.009*	5.086**	5.927***	9.919***	9.657***	12.571***	12.543***	9.299***	10.707***
	(1.86)	(2.30)	(2.60)	(4.14)	(4.03)	(4.60)	(4.59)	(3.09)	(3.51)
MSinceAcqui	0.012***	0.029***	0.029***	0.013***	0.013***	0.017***	0.017***	0.019***	0.020***
	(15.81)	(19.58)	(18.71)	(6.79)	(7.15)	(7.89)	(7.89)	(7.85)	(7.90)
PatentsCount	0.062***	0.102***	0.101***	0.037***	0.032***	0.032***	0.032***	0.033***	0.033***
	(8.14)	(11.20)	(10.95)	(4.37)	(3.85)	(3.45)	(3.46)	(3.45)	(3.42)
FirstFiling		0.190***	0.191***	-0.018	-0.008	-0.018	-0.018	-0.023	-0.026
		(14.05)	(13.21)	(-0.93)	(-0.41)	(-0.83)	(-0.83)	(-0.93)	(-1.02)
US			-0.262***	-0.251***	-0.250***	-0.299**	-0.149	-0.223	-0.730
			(-2.98)	(-2.76)	(-2.74)	(-2.30)	(-0.37)	(-0.50)	(-1.28)
MSinceLastFil				-0.031***	-0.030***	-0.031***	-0.031***	-0.038***	-0.039***
				(-14.56)	(-14.10)	(-13.13)	(-13.12)	(-13.54)	(-13.71)
FirstAuthor					0.272***	0.314***	0.314***	0.244***	0.235***
					(3.87)	(4.02)	(4.02)	(2.87)	(2.70)
Incorp^T						0.013	0.013	-0.012	-0.018
						(1.36)	(1.35)	(-1.15)	(-1.60)
US^T							-0.152	-0.004	0.476
							(-0.40)	(-0.01)	(0.89)
FirmSize^T								-0.604***	-0.601***
								(-7.45)	(-5.02)
Funding^T									0.029
									(0.23)
Constant	-1.668***	-386.253***	-386.660***	35.282	14.925	9.500	9.573	69.737	86.885*
	(-13.02)	(-14.11)	(-13.25)	(0.90)	(0.38)	(0.21)	(0.21)	(1.37)	(1.66)
Observations	4880	4880	4720	4720	4720	3967	3967	3509	3435
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

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

Table B.18: Inventors innovating for their acquirer, MarketValue definition (b.3)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Talent									
Core=1	0.447***	0.413***	0.380***	0.342***	0.338***	0.553***	0.553***	0.529***	0.715***
	(4.79)	(4.17)	(3.74)	(3.28)	(3.24)	(3.95)	(3.94)	(2.91)	(3.68)
MarketVal_ter	-0.662	-1.786***	-1.898***	-1.124**	-0.937*	-1.900***	-1.900***	-2.609***	-2.733***
	(-1.36)	(-3.48)	(-3.68)	(-2.19)	(-1.81)	(-3.13)	(-3.13)	(-4.05)	(-4.19)
Core=1 × MarketVal_ter	0.524	0.922	1.137*	0.953*	0.852	1.967***	1.967***	2.446***	2.618***
	(0.94)	(1.53)	(1.87)	(1.66)	(1.48)	(2.97)	(2.97)	(3.48)	(3.69)
MSinceAcqui	0.009***	0.027***	0.026***	0.009***	0.010***	0.014***	0.014***	0.017***	0.018***
	(11.29)	(18.08)	(17.18)	(4.87)	(5.18)	(6.38)	(6.38)	(6.44)	(6.45)
PatentsCount	0.063***	0.109***	0.106***	0.029***	0.023**	0.020**	0.020**	0.021**	0.020**
	(7.38)	(10.76)	(10.21)	(3.05)	(2.48)	(2.02)	(2.02)	(2.18)	(2.07)
FirstFiling		0.210***	0.205***	-0.008	0.001	-0.007	-0.007	-0.008	-0.008
		(15.11)	(14.01)	(-0.40)	(0.04)	(-0.32)	(-0.32)	(-0.31)	(-0.30)
US			-0.207**	-0.164	-0.157	-0.273**	-0.273	-0.086	0.092
			(-2.13)	(-1.64)	(-1.56)	(-2.06)	(-0.57)	(-0.15)	(0.13)
MSinceLastFil				-0.029***	-0.028***	-0.030***	-0.030***	-0.033***	-0.035***
				(-13.75)	(-13.40)	(-12.74)	(-12.73)	(-12.38)	(-12.51)
FirstAuthor					0.303***	0.318***	0.318***	0.245***	0.228**
					(4.03)	(3.82)	(3.82)	(2.69)	(2.45)
Incorp^T						0.022**	0.022**	0.011	-0.001
						(2.02)	(2.01)	(0.85)	(-0.04)
US^T							0.000	-0.150	-0.169
							(0.00)	(-0.27)	(-0.25)
FirmSize^T								-0.565***	-0.441***
								(-6.15)	(-3.31)
Funding^T									-0.115
									(-0.84)
Constant	-1.125***	-425.164***	-415.995***	16.344	-2.058	-30.108	-30.108	-6.191	17.849
	(-7.14)	(-15.15)	(-14.04)	(0.40)	(-0.05)	(-0.61)	(-0.61)	(-0.11)	(0.31)
Observations	4158	4158	4009	4009	4009	3383	3383	2955	2881
BT dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

t statistics in parentheses

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

B.11 Probit estimates of the selection equation

Table B.19: Heckman selection equation (2.4)

	Talent
MarketVal	-1.413***
	(0.349)
Core_Before = 1	0.305***
	(0.071)
Core_Before=1 × MarketVal	2.241***
	(0.416)
MSinceAcqui	0.004***
	(0.001)
FirstFiling	-0.037***
	(0.013)
US	0.036
	(0.068)
FirstAuthorSh	0.026
	(0.067)
MSinceLastFiling	-0.011***
	(0.002)
Constant	73.462***
	(25.565)
Observations	2,515
BT dummies	Yes
Prob > $\chi 2$	0,00
Pseudo R^2	0,10

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The z-test for the significance of the coefficient of $MSinceLastFiling$ is associated with a p-value < 0.01; we conclude that this is a relevant instrument.

B.12 Selection model – Alternative Market Value Indices

Table B.20: Heckman two-step parameters, MarketValue definition (a)

	(1) first	(2) select	(3) /mills
MarketVal_bar	0.056	-1.514***	
	(0.173)	(0.498)	
Core_Before = 1	0.760***	-1.046***	
	(0.073)	(0.219)	
Core_Before=1 × MarketVal	-0.345	4.269***	
	(0.246)	(0.686)	
MSinceAcqui	-0.000	0.004***	
	(0.000)	(0.001)	
FirstFiling	0.004	-0.045***	
	(0.003)	(0.012)	
US	-0.085***	0.038	
	(0.019)	(0.068)	
FirstAuthorSh	0.041**	0.022	
	(0.019)	(0.067)	
MSinceLastFiling		-0.011***	
		(0.002)	
lambda			-0.036
			(0.064)
Constant	-8.025	90.070***	
	(5.829)	(24.611)	
Observations	2,515	2,515	2,515
BT dummies	Yes	Yes	Yes

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table B.21: Heckman two-step parameters, MarketValue definition (b.2)

	(1) first	(2) select	(3) /mills
MarketVal_bis	0.248	-3.262**	
	(0.453)	(1.587)	
Core_Before = 1	0.634***	0.412***	
	(0.023)	(0.075)	
Core_Before=1 × MarketVal	-2.003***	16.110***	
	(0.696)	(2.283)	
MSinceAcqui	-0.000	0.004***	
	(0.000)	(0.001)	
FirstFiling	0.005*	-0.046***	
	(0.003)	(0.012)	
US	-0.077***	-0.010	
	(0.019)	(0.069)	
FirstAuthorSh	0.039**	0.025	
	(0.019)	(0.068)	
MSinceLastFiling		-0.011***	
		(0.002)	
lambda			-0.087
			(0.061)
Constant	-10.668*	91.533***	
	(5.981)	(24.807)	
Observations	2,515	2,515	2,515
BT dummies	Yes	Yes	Yes

Standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Table B.22: Heckman two-step parameters, MarketValue definition (b.3)

	(1) first	(2) select	(3) /mills
MarketVal_ter	0.032	-0.570*	
	(0.075)	(0.334)	
Core_Before = 1	0.647***	0.238***	
	(0.022)	(0.082)	
Core_Before=1 × MarketVal	-0.188**	0.908**	
	(0.089)	(0.362)	
MSinceAcqui	0.000	0.001	
	(0.000)	(0.001)	
FirstFiling	0.003	-0.033***	
	(0.003)	(0.013)	
US	-0.121***	0.092	
	(0.019)	(0.079)	
FirstAuthorSh	0.041**	0.063	
	(0.019)	(0.074)	
MSinceLastFiling		-0.008***	
		(0.002)	
lambda			0.017
			(0.085)
Constant	-6.623	66.042***	
	(5.596)	(25.558)	
Observations	1,965	1,965	1,965
BT dummies	Yes	Yes	Yes

Standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Appendix C

Table C.1: Bureau Van Dijk's peer groups to which Big Tech targets belong

Peer Group Name	PeerGroupSize	TargetCount	TargetShare
Computer programming activities	76511	117	.002
Other software publishing	41458	54	.001
Other information technology and computer service activities	84145	36	0
Other professional, scientific and technical activities nec	259502	19	0
Retail sale via mail order houses or via Internet	18977	15	.001
Other telecommunications activities	36156	11	0
Computer consultancy activities	21166	11	.001
Service activities incidental to water transportation	15484	10	.001
Manufacture of electronic components	50222	9	0
Manufacture of computers and peripheral equipment	13044	9	.001
Engineering activities and related technical consultancy	175964	5	0
Data processing, hosting and related activities	29887	5	0
Advertising agencies	82905	5	0
Wholesale of computers, computer peripheral equipment and software	47547	4	0
Other research and experimental development on natural sciences and engineering	32895	4	0
Other financial service activities, except insurance and pension funding nec	45773	4	0
Other business support service activities nec	209099	4	0
Motion picture, video and television programme post-production activities	5040	4	.001
Retail sale of computers, peripheral units and software in specialised stores	53034	3	0
Non-specialised wholesale trade	298724	3	0
Manufacture of instruments and appliances for measuring, testing and navigation	21472	3	0
Manufacture of communication equipment	13254	3	0
Wired telecommunications activities	11370	2	0
Specialised design activities	16708	2	0
Retail sale of games and toys in specialised stores	3725	2	.001
Restaurants and mobile food service activities	212558	2	0
Other retail sale not in stores, stalls or markets	48526	2	0
Manufacture of optical instruments and photographic equipment	5171	2	0
Industrial companies	642	2	.003
Business and other management consultancy activities	184239	2	0
Artistic creation	7885	2	0
Wireless telecommunications activities	5056	1	0
Wholesale of electrical household appliances	70351	1	0
Wholesale of beverages	46863	1	0
Web portals	10299	1	0
Travel agency activities	41417	1	0
Sound recording and music publishing activities	4685	1	0
Security systems service activities	11329	1	0
Retail sale via stalls and markets of food, beverages and tobacco products	11557	1	0
Retail sale of newspapers and stationery in specialised stores	5530	1	0
Retail sale of cosmetic and toilet articles in specialised stores	13644	1	0
Retail sale of clothing in specialised stores	170841	1	0
Retail sale of books in specialised stores	42470	1	0
Retail sale of audio and video equipment in specialised stores	43391	1	0
Retail sale in non-specialised stores with food, beverages or tobacco predominating	234929	1	0
Reproduction of recorded media	1668	1	.001
Repair of electrical equipment	7091	1	0
Radio broadcasting	13141	1	0
Publishing of computer games	397	1	.003
Primary education	58746	1	0
Photographic activities	14387	1	0
Other retail sale of new goods in specialised stores	290929	1	0
Other reservation service and related activities	38460	1	0
Other publishing activities	4686	1	0
Other information service activities nec	22071	1	0
Other human health activities	119835	1	0
Other education nec	43212	1	0
Motion picture, video and television programme production activities	27004	1	0
Motion picture, video and television programme distribution activities	3629	1	0
Manufacture of other general-purpose machinery nec	38938	1	0
Manufacture of other electrical equipment	24688	1	0
Manufacture of motor vehicles	4280	1	0
Manufacture of consumer electronics	4173	1	0
Manufacture of air and spacecraft and related machinery	4875	1	0
Legal activities	83934	1	0
Hotels and similar accommodation	91661	1	0
Fitness facilities	21586	1	0
Dispensing chemist in specialised stores	131204	1	0
Agents specialised in the sale of other particular products	71698	1	0
Activities of holding companies	213986	1	0

C.1 Big Tech-acquired patents portfolios

The tables below compare the descriptive statistics of all Big Tech-acquired patents (see Table C.2) with those matched with Arts, Hou, and Gomez (2021)'s database (see Table C.3).

Table C.2: Big Tech-acquired patents portfolios

		Portfolio size (patents #)		Patent age at acquisition (y)		Filing date	Grant date
	Count	Mean	SD	Mean	SD	[min,max]	[min,max]
AMZN	27	22.07	64.62	3.31	2.59	1995m6, 2020m6	1998m8, 2022m6
APPL	52	14.21	19.72	4.00	2.59	1988m6, 2019m11	1990m5, 2022m5
FCBK	18	7.56	17.68	4.31	4.37	1998m11, 2019m1	2001m1, 2020m6
GOOG	67	30.98	143.16	3.90	2.14	1995m6, 2020m9	1996m12, 2022m6
MSFT	88	16.52	51.88	3.86	2.79	1987m5, 2020m4	1989m2, 2022m6
TOTAL	252	19.84	83.12	3.87	2.71	1987m5, 2020m9	1989m2, 2022m6

Table C.3: Big Tech-acquired patents portfolios with textual data

		Portfolio size (patents #)		Patent age at acquisition (y)		Share of patents with new keywords pairs		Filing date	Grant date
	Count	Mean	SD	Mean	SD	Mean	SD	[min,max]	[min,max]
AMZN	27	22.07	64.62	3.60	2.61	.83	.26	1995m6 , 2017m9	1998m8 , 2018m5
APPL	47	15.28	20.44	4.43	2.68	.87	.21	1988m6 , 2017m6	1990m5 , 2018m5
FCBK	14	8.79	19.95	4.88	4.87	.93	.14	1998m11 , 2015m5	2001m1 , 2018m3
GOOG	65	31.88	145.28	4.11	2.27	.90	.22	1995m6 , 2017m5	1996m12 , 2018m5
MSFT	80	17.9	54.24	4.36	3.11	.86	.26	1987m5 , 2017m7	1989m2 , 2018m5
TOTAL	233	21.21	86.31	4.25	2.89	.87	.23	1987m5, 2017m9	1989m2, 2018m5

C.2 Disruption metric - Extension

In order to directly integrate the notion of future impact in the disruption metric, we propose to weight the number of new keywords combinations in a patent p by an index comprised between 0 and 1 capturing the number of forward citations received by this patent:

$$Disrupt_p^* = \widetilde{FwdCit}_p * NewKeywords_p$$

where the citations-based weights have been normalized such as to be comprised between 0 and 1 by min/max scaling:

$$\widetilde{FwdCit} = (FwdCit - min(FwdCit)) / (max(FwdCit) - min(FwdCit)).$$

The number of citations received by a patent ($FwdCit_p$) is a stock that builds over time, so patents published at different times cannot be compared. To overcome this problem, we propose to consider the number of forward citations received by a patent *over a period of five years* after its publication date. This means that this variable can only be used for patents published at least five years before the end of our study period (in July 2022), so in July 2017. Because publication typically occurs around 18 months after the filing date (Squicciarini, Dernis, and Criscuolo 2013), we restrict this alternative specification of our disruption metric to patents filed before January 2016.

C.3 Big Tech targets' product markets

CATEGORY	STATISTA MARKET	CRUNCHBASE MARKET
Entertainment	Digital Video	Digital Video, Video, Video Streaming
	ePublishing	ePublishing, Publishing, Audiobooks
	Digital Music	Digital Music, Music, Music Streaming
	Extended Reality	Extended Reality, Augmented Reality, 3D Technology, Virtual Reality
	Game Development	Game Development, Video Games, Console Games, PC Games

Connected Devices	Smart Home	Smart Home
	Internet of Things	Internet of Things
Apps	Digital Health	Digital Health, Fitness, Wellness, Personal Health
	Education	Education Apps, Education
Hardware	Hardware	Hardware, Mobile Devices
	Semiconductor	Semiconductor
Social Interactions	Mobile Messaging	Mobile Messaging, Messaging
	Social Network	Social Network, Photo Sharing, Social Media
Imaging	Mapping	Mapping, Navigation
	Digital Imaging Technology	Digital Imaging Technology
Artificial Intelligence	Robotics	Robotics, Autonomous Vehicles, Neuroscience
	Image Recognition	Image Recognition, Computer Vision, Facial Recognition

	Speech-Based NLP	Speech-Based NLP, Speech Recognition
	Generative AI	Generative AI, Text-to-Speech
	Business Intelligence	Business Intelligence, Data Mining, Marketing Analytics
	Language Translation NLP	Language Translation NLP
	Text-based NLP	Text-based NLP, Text Analytics
	Predictive Analytics	Predictive Analytics, Virtual Agent
Consumers	eCommerce	E-Commerce, Shopping, Retail, Comparison Shopping Engine
	In App Advertising	In App Advertising
	CRM	CRM, Mobile Payment, Consumer, Advertising
Enterprise Softwares	ERP	ERP, eCommerce Software
	SCM	SCM, Logistics
	Collaboration Software	Collaboration Software, Video Chat, Video Conferencing, Collaboration, File Sharing
	Content Management Software	Content Management Software
	Creative Software	Creative Software

	Construction and Design Software	Construction and Design Software, Product Design
Development	Application Development Software	Application Development Software, Developer APIs, Web Development, Developer Tools
	Application Outsourcing	Application Outsourcing
	PaaS	PaaS, Productivity Tools, Developer Platform
Network Systems	Virtualization	Virtualization
	Cloud Storage	Cloud Storage, Data Storage, Cloud Computing, Data Recovery
	Cyber Security	Cyber Security, Network Security, Cloud Security
	System Infrastructure Software	System Infrastructure Software
Others	Lighting	Lighting
	Energy	Energy
	Medical Device	Medical Device

This table was constructed by webscraping from the Crunchbase website the products in which Big Tech targets are active, and matching them with Statista's market data.

C.4 Distribution of the dependent variable

Figure C.1: New keywords combinations

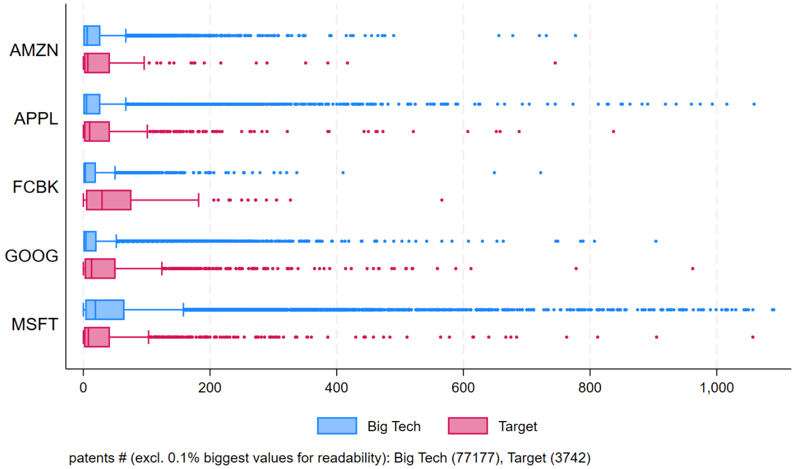

Figure C.2: New keywords combinations, excluding outliers

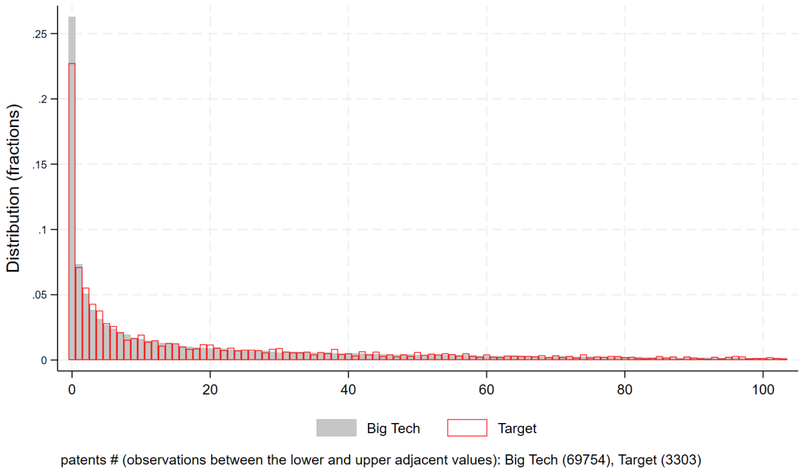

C.5 Average disruption levels – Comparing Big Tech and its targets

Table C.5: New keywords combinations

	(1) Big Tech	(2) Targets	(3) Mean Difference
AMZN	23.61	84.55	-60.94
	(.)	(149.91)	
APPL	28.90	98.91	-70.01
	(.)	(184.94)	
FCBK	18.94	63.71	-44.77
	(.)	(76.98)	
GOOG	21.06	94.65	-73.59
	(.)	(189.36)	
MSFT	56.82	84.87	-28.05
	(.)	(218.83)	
Total	29.86	89.12	-59.26
	(15.52)	(189.49)	
N	5	233	238

Table C.6: New keywords combinations, citations-weighted

	(1) Big Tech	(2) Targets	(3) Mean Difference
AMZN	0.18	0.37	-0.19
	(.)	(0.70)	
APPL	0.29	1.04	-0.75
	(.)	(2.28)	
FCBK	0.18	0.34	-0.16
	(.)	(0.50)	
GOOG	0.19	1.22	-1.02
	(.)	(3.51)	
MSFT	0.37	0.57	-0.19
	(.)	(1.00)	
Total	0.24	0.81	-0.57
	(0.09)	(2.23)	
N	5	229	234

Table C.7: New keywords combinations, excl. 0

	(1) Big Tech	(2) Targets	(3) Mean Difference
AMZN	32.35	93.48	-61.13
	(.)	(150.47)	
APPL	40.36	107.85	-67.49
	(.)	(187.42)	
FCBK	28.86	65.95	-37.10
	(.)	(77.83)	
GOOG	31.11	100.88	-69.76
	(.)	(190.89)	
MSFT	67.56	92.60	-25.04
	(.)	(223.07)	
Total	40.05	96.48	-56.43
	(15.98)	(191.78)	
N	5	225	230

Table C.8: New keywords combinations, citations-weighted and excl. 0

	(1) Big Tech	(2) Targets	(3) Mean Difference
AMZN	0.26	0.43	-0.17
	(.)	(0.72)	
APPL	0.42	1.16	-0.74
	(.)	(2.43)	
FCBK	0.28	0.36	-0.08
	(.)	(0.50)	
GOOG	0.30	1.29	-0.99
	(.)	(3.56)	
MSFT	0.46	0.62	-0.16
	(.)	(1.01)	
Total	0.34	0.89	-0.54
	(0.09)	(2.30)	
N	5	221	226

C.6 Model 3.1 – Robustness checks

Table C.9: Disruption of Internally developed vs Acquired top patents, citations-weighted

	(1)	(2)	(3)	(4)	(5)
Disrupt*(*index*)					
Acquired=1	1.021***	0.822***	0.923***	0.979***	0.790***
	(23.63)	(18.80)	(21.45)	(23.02)	(17.40)
MSFT=1	-0.119***	-0.117***	-0.045**	0.124***	0.141***
	(-5.71)	(-5.59)	(-2.11)	(5.78)	(6.47)
Acquired=1 × MSFT=1	-0.600***	-0.442***	-0.612***	-0.803***	-0.616***
	(-7.11)	(-5.18)	(-7.06)	(-9.44)	(-7.13)
Patent Scope		0.302***	0.209***	0.201***	0.141***
		(37.72)	(25.91)	(25.21)	(15.34)
Family Size			0.133***	0.119***	0.118***
			(50.52)	(45.35)	(44.58)
Number of Claims				0.035***	0.035***
				(44.19)	(44.32)
Constant	-3.641	-3.943	-3.983	-4.348	-4.617
	(-0.57)	(-0.62)	(-0.63)	(-0.69)	(-0.73)
Observations	73482	73482	73482	73480	73476
Pseudo R^2	0.056	0.071	0.098	0.120	0.130
Year dummies	Yes	Yes	Yes	Yes	Yes
CPC dummies	No	No	No	No	Yes

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C.10: Disruption of Internally developed vs Acquired top patents, excl. 0

	(1)	(2)	(3)	(4)	(5)
Disrupt (excl. 0)					
Acquired=1	0.491***	0.478***	0.476***	0.512***	0.486***
	(16.29)	(15.74)	(15.67)	(17.05)	(16.04)
MSFT=1	0.382***	0.382***	0.384***	0.483***	0.485***
	(32.52)	(32.51)	(32.62)	(40.55)	(40.13)
Acquired=1 × MSFT=1	-0.410***	-0.400***	-0.407***	-0.504***	-0.467***
	(-7.53)	(-7.33)	(-7.46)	(-9.32)	(-8.60)
Patent Scope		0.025***	0.020***	0.016***	0.013**
		(5.03)	(4.00)	(3.17)	(2.43)
Family Size			0.011***	0.003*	0.004**
			(6.37)	(1.83)	(2.16)
Number of Claims				0.023***	0.023***
				(41.71)	(41.55)
Constant	3.367***	3.342***	3.336***	3.092**	3.165**
	(2.64)	(2.62)	(2.61)	(2.45)	(2.52)
Observations	61642	61420	61420	61418	61414
Pseudo R^2	0.014	0.014	0.014	0.017	0.018
Year dummies	Yes	Yes	Yes	Yes	Yes
CPC dummies	No	No	No	No	Yes

t statistics in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Bibliography

Abrams, David S, Ufuk Akcigit, and Jillian Grennan (2013). *Patent value and citations: Creative destruction or strategic disruption?* Tech. rep. National Bureau of Economic Research.

ACCC (2019). "Digital platforms inquiry. Final Report". *Australian Competition and Consumer Commission.*

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, et al. (2018). "Innovation, reallocation, and growth". *American Economic Review* 108.11, pp. 3450–3491.

Acemoglu, Daron, Ufuk Akcigit, and Murat Alp Celik (2014). *Young, restless and creative: Openness to disruption and creative innovations*. Tech. rep. National Bureau of Economic Research.

Adner, Ron (2002). "When are technologies disruptive? A demand-based view of the emergence of competition". *Strategic management journal* 23.8, pp. 667–688.

Affeldt, Pauline and Reinhold Kesler (2021). "Competitors' Reactions to Big Tech Acquisitions: Evidence from Mobile Apps". *DIW Discussion Papers* 1987.

Aghion, Philippe et al. (2005). "Competition and innovation: An inverted-U relationship". *The quarterly journal of economics* 120.2, pp. 701–728.

Ahuja, Gautam and Riitta Katila (2001). “Technological acquisitions and the innovation performance of acquiring firms: A longitudinal study”. *Strategic management journal* 22.3, pp. 197–220.

Ajiferuke, Isola and Felix Famoye (2015). “Modelling count response variables in informetric studies: Comparison among count, linear, and lognormal regression models”. *Journal of Informetrics* 9.3, pp. 499–513.

Akcigit, Ufuk and Nathan Goldschlag (2023). *Where have all the" creative talents" gone? Employment dynamics of us inventors*. Tech. rep. National Bureau of Economic Research.

Akcigit, Ufuk and William R Kerr (2018). “Growth through heterogeneous innovations”. *Journal of Political Economy* 126.4, pp. 1374–1443.

Allison, John R and Mark A Lemley (1998). “Empirical evidence on the validity of litigated patents”. *AIPLA QJ* 26, p. 185.

Anderson, Brian (2017). *Selection Models and Weak Instruments*. Accessed on February 28, 2024. URL: `https://info.umkc.edu/drbanderson/selection-models-and-weak-instruments/#`.

Antonio, Jerome L and Dominik K Kanbach (2023). “Contextual factors of disruptive innovation: A systematic review and framework”. *Technological Forecasting and Social Change* 188, p. 122274.

Argentesi, Elena et al. (2021). “Merger policy in digital markets: an ex post assessment”. *Journal of Competition Law & Economics* 17.1, pp. 95–140.

Arrow, Kenneth Joseph (1972). *Economic welfare and the allocation of resources for invention*. Springer.

Arts, Sam, Bruno Cassiman, and Jianan Hou (2021). *Technology differentiation, product market rivalry, and M&A transactions*. Tech. rep. Working paper.

Arts, Sam, Jianan Hou, and Juan Carlos Gomez (2021). "Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measures". *Research Policy* 50.2, p. 104144.

Arundel, Anthony (2001). "The relative effectiveness of patents and secrecy for appropriation". *Research policy* 30.4, pp. 611–624.

Asdrubali, Pierfederico and Simone Signore (2015). *The economic impact of EU guarantees on credit to SMEs–Evidence from CESEE countries*. Tech. rep. EIF Working Paper.

Austin, Peter C (2011). "An introduction to propensity score methods for reducing the effects of confounding in observational studies". *Multivariate behavioral research* 46.3, pp. 399–424.

Bar-Isaac, Heski, Justin Pappas Johnson, and Volker Nocke (2023). *Acquihiring for Monopsony Power*. Centre for Economic Policy Research.

Barefoot, Kevin et al. (2018). "Defining and measuring the digital economy". *US Department of Commerce Bureau of Economic Analysis, Washington, DC* 15, p. 210.

Belleflamme, Paul and Martin Peitz (2015). *Industrial organization: markets and strategies*. Cambridge University Press.

Ben-Slimane, Karim, Cédric Diridollou, and Karim Hamadache (2020). "The legitimation strategies of early stage disruptive innovation". *Technological Forecasting and Social Change* 158, p. 120161.

Benkert, Jean-Michel, Igor Letina, and Shuo Liu (2023). "Startup Acquisitions: Acquihires and Talent Hoarding". *arXiv preprint arXiv:2308.10046*.

Block, Joern H, Christian O Fisch, and Mirjam Van Praag (2017). "The Schumpeterian entrepreneur: a review of the empirical evidence on the antecedents, behaviour and consequences of innovative entrepreneurship". *Industry and innovation* 24.1, pp. 61–95.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2024). "Revisiting event-study designs: robust and efficient estimation". *Review of Economic Studies*, rdae007.

Bourreau, Marc, Bruno Jullien, and Yassine Lefouili (2021). "Mergers and demand-enhancing innovation". *CEPR Discussion Paper No. DP16031*.

Boyacıoğlu, Beril, Mahmut N Özdemir, and Samina Karim (2024). "Acqui-hires: Redeployment and retention of human capital post-acquisition". *Strategic Management Journal* 45.2, pp. 205–237.

Bresnahan, Timothy F, Shane Greenstein, and Rebecca M Henderson (2011). "Schumpeterian competition and diseconomies of scope: illustrations from the histories of Microsoft and IBM". *The rate and direction of inventive activity revisited*. University of Chicago Press, pp. 203–271.

Bro, Rasmus and Age K Smilde (2014). "Principal component analysis". *Analytical methods* 6.9, pp. 2812–2831.

Bryan, Kevin A and Erik Hovenkamp (2020a). "Antitrust limits on startup acquisitions". *Review of Industrial Organization* 56.4, pp. 615–636.

— (2020b). "Startup acquisitions, error costs, and antitrust policy". *The University of Chicago Law Review* 87.2, pp. 331–356.

Bryan, Kevin A and Yasin Ozcan (2021). "The impact of open access mandates on invention". *Review of Economics and Statistics* 103.5, pp. 954–967.

Bryan, Kevin A, Yasin Ozcan, and Bhaven Sampat (2020). "In-text patent citations: A user's guide". *Research Policy* 49.4, p. 103946.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2019). "Artificial intelligence and the modern productivity paradox". *The economics of artificial intelligence: An agenda* 23, pp. 23–57.

BWB and BKartA (2018). "Guidance on Transaction Value Thresholds for Mandatory Pre-merger Notification". Section 35 (1a) of the German Competition Act and Section 9 (4) of the Austrian Cartel Act 2005.

Cabral, Luis (2017). *Introduction to industrial organization*. MIT press.

— (2018). "Standing on the shoulders of dwarfs: Dominant firms and innovation incentives". *Available at SSRN 3235598*.

— (2021). "Merger policy in digital industries". *Information Economics and Policy* 54, p. 100866.

— (2023). *Big Tech Acquisitions*. Centre for Economic Policy Research.

Caffarra, Cristina, Gregory Crawford, and Tommaso Valletti (2020). "How tech rolls': Potential competition and 'reverse'killer acquisitions". *Antitrust Chronicle* 2.2, pp. 1–9.

Campbell, Benjamin A et al. (2012). "Who leaves, where to, and why worry? Employee mobility, entrepreneurship and effects on source firm performance". *Strategic Management Journal* 33.1, pp. 65–87.

Cassiman, Bruno, Massimo G Colombo, et al. (2005). "The impact of M&A on the R&D process: An empirical analysis of the role of technological-and market-relatedness". *Research policy* 34.2, pp. 195–220.

Cassiman, Bruno and Reinhilde Veugelers (2006). "In search of complementarity in innovation strategy: Internal R&D and external knowledge acquisition". *Management science* 52.1, pp. 68–82.

Cattell, Raymond B (1966). "The scree test for the number of factors". *Multivariate behavioral research* 1.2, pp. 245–276.

Caves, Richard E (1989). "Mergers, takeovers, and economic efficiency: foresight vs. hindsight". *International Journal of Industrial Organization* 7.1, pp. 151–174.

Certo, S Trevis et al. (2016). "Sample selection bias and Heckman models in strategic management research". *Strategic Management Journal* 37.13, pp. 2639–2657.

Cheng, Yu et al. (2017). "Forecasting of potential impacts of disruptive technology in promising technological areas: Elaborating the SIRS epidemic model in RFID technology". *Technological Forecasting and Social Change* 117, pp. 170–183.

Christensen, Clayton M (1997). *The innovator's dilemma: when new technologies cause great firms to fail*. Harvard Business Review Press.

— (2013). *The innovator's dilemma: when new technologies cause great firms to fail*. Harvard Business Review Press.

Cloodt, Myriam, John Hagedoorn, and Hans Van Kranenburg (2006). "Mergers and acquisitions: Their effect on the innovative performance of companies in high-tech industries". *Research policy* 35.5, pp. 642–654.

Clougherty, Joseph A, Tomaso Duso, and Johannes Muck (2016). "Correcting for self-selection based endogeneity in management research: Review, recommendations and simulations". *Organizational Research Methods* 19.2, pp. 286–347.

CMA (2021). "CMA directs Facebook to sell Giphy." URL: https://www.gov.uk/government/news/cma-directs-facebook-to-sell-giphy.

Cohen, Wesley M and Richard C Levin (1989). "Empirical studies of innovation and market structure". *Handbook of industrial organization* 2, pp. 1059–1107.

Cook, Tim (2019). "Berkshire Hathaway's annual shareholder meeting." *Interview by Becky Quick on CNBC*.

Crémer, Jacques (Oct. 2020). "Roundtable "Conference on Digital Platforms: Opportunities and Challenges?" Online", p. 8. URL: https://www.tse-fr.eu/sites/default/files/TSE/documents/Rapports/tse-digital-annual-report-2020.pdf#page=4.

Crémer, Jacques, Yves-Alexandre de Montjoye, and Heike Schweitzer (2019). "Competition policy for the digital era. Final report". *European Commission, DG Comp: Brussels*.

Criscuolo, Chiara (2019). "Changes in productivity and industry dynamics in the digital transition: cross-country evidence from firm-level data". *CLPD* 5, p. 21.

Crouzet, Nicolas and Janice C Eberly (2019). *Understanding weak capital investment: The role of market concentration and intangibles.* Tech. rep. National Bureau of Economic Research.

Cunningham, Colleen, Florian Ederer, and Song Ma (Mar. 2021). "Killer Acquisitions". en. *Journal of Political Economy* 129.3, pp. 649–702. ISSN: 0022-3808, 1537-534X.

De Chaisemartin, Clément and Xavier d'Haultfoeuille (2024). "Difference-in-differences estimators of intertemporal treatment effects". *Review of Economics and Statistics*, pp. 1–45.

De Coninck, Raphaël and Christoph von Muellern (2023). "Big Tech Acquisitions and Innovation Incentives". *Network Law Review.*

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger (2020). "The rise of market power and the macroeconomic implications". *The Quarterly Journal of Economics* 135.2, pp. 561–644.

Dehejia, Rajeev H and Sadek Wahba (2002). "Propensity score-matching methods for nonexperimental causal studies". *Review of Economics and statistics* 84.1, pp. 151–161.

Denicolò, Vincenzo and Michele Polo (2018). "Duplicative research, mergers and innovation". *Economics Letters* 166, pp. 56–59.

"Digital Markets Act" (2020). *Official Journal of the European Union* 2020/0374 (COD).

Dijk, Esmée, José L Moraga-González, and Evgenia Motchenkova (2021). "How do start-up acquisitions affect the direction of innovation?" *The Journal of Industrial Economics.*

Dijk, Esmée, José Luis Moraga-González, and Evgenia Motchenkova (2021). "How do start-up acquisitions affect the direction of innovation?" *The Journal of Industrial Economics.*

Doan, Thanh and Franco Mariuzzo (2022). "Digital Platform Mergers and Innovation: Evidence From Cloud Computing Market".

"Dow/DuPont" (2017). *Case M.7932, Commission decision of 27 March 2017.*

EC (2021). "Communication from the Commission Guidance on the application of the referral mechanism set out in Article 22 of the Merger Regulation, art. 2.2(19)". *Official Journal of the European Union* 2021/C 113/01.

Efrati, A. (2015). *Google's war against apps.* URL: `https://www.theinformation.com/articles/google-s-war-against-apps`.

Eisfeld, Luise (2022). *Entry and acquisitions in software markets.* Tech. rep. Working Paper, Toulouse School of Economics.

Esteva Mosso, Carles (2018). "Innovation in EU Merger Control". *European Commission, DG Comp: Brussels.*

Etro, Federico (2019). "Mergers of complements and entry in innovative industries". *International Journal of Industrial Organization* 65, pp. 302–326.

EU Scoreboard, European Commission (2021). "The 2021 EU Industrial R&D Investment Scoreboard".

Ezrachi, Ariel and Maurice E Stucke (2022). *How Big-Tech Barons Smash Innovation—and How to Strike Back.* HarperCollins.

Federico, Giulio, Gregor Langus, and Tommaso Valletti (2018). "Horizontal mergers and product innovation". *International Journal of Industrial Organization* 59, pp. 1–23.

Federico, Giulio, Fiona Scott Morton, and Carl Shapiro (2020). "Antitrust and innovation: Welcoming and protecting disruption". *Innovation Policy and the Economy* 20.1, pp. 125–190.

Fons-Rosen, Christian, Pau Roldan-Blanco, and Tom Schmitz (2021). "The Effects of Startup Acquisitions on Innovation and Economic Growth". *Max Planck Institute for Innovation & Competition Research Paper* 23-02.

Fumagalli, Chiara, Massimo Motta, and Emanuele Tarantino (2020). "Shelving or developing? The acquisition of potential competitors under financial constraints". *CEPR Discussion Paper No. DP15113*.

Funk, Russell J and Jason Owen-Smith (2017). "A dynamic network measure of technological change". *Management science* 63.3, pp. 791–817.

Furman, Jason et al. (2019). "Unlocking digital competition: Report of the digital competition expert panel". *UK government publication, HM Treasury* 27.

Galloway, Scott (2018). *The four: the hidden DNA of Amazon, Apple, Facebook, and Google*. Penguin.

Gambardella, Alfonso, Dietmar Harhoff, and Bart Verspagen (2008). "The value of European patents". *European Management Review* 5.2, pp. 69–84.

Gans, Joshua (2016). *The disruption dilemma*. MIT press.

Gautier, Axel and Joe Lamesch (2021). "Mergers in the digital economy". *Information Economics and Policy* 54, p. 100890.

Gdula, Olga Ewa (2021). "Essays on empirical econometrics". PhD thesis. University of Oxford.

Govindarajan, Vijay and Praveen K Kopalle (2006). "Disruptiveness of innovations: measurement and an assessment of reliability and validity". *Strategic management journal* 27.2, pp. 189–199.

Grant, Robert M (1996). "Toward a knowledge-based theory of the firm". *Strategic management journal* 17.S2, pp. 109–122.

Gugler, Klaus, Florian Szücs, and Ulrich Wohak (2023). "Start-up Acquisitions, Venture Capital and Innovation: A Comparative Study of Google, Apple, Facebook, Amazon and Microsoft". *WU Vienna University of Economics and Business. Working Paper Series No. 340*.

Hahn, Jinyong and Jerry Hausman (2003). "Weak instruments: Diagnosis and cures in empirical econometrics". *American Economic Review* 93.2, pp. 118–125.

Harhoff, Dietmar and Stefan Wagner (2009). "The duration of patent examination at the European Patent Office". *Management Science* 55.12, pp. 1969–1984.

Hashmi, Aamir Rafique and Johannes Van Biesebroeck (2016). "The relationship between market structure and innovation in industry equilibrium: a case study of the global automobile industry". *Review of Economics and Statistics* 98.1, pp. 192–208.

Haucap, Justus, Alexander Rasch, and Joel Stiebale (2019). "How mergers affect innovation: Theory and evidence". *International Journal of Industrial Organization* 63, pp. 283–325.

Heckman, James J (1979). "Sample selection bias as a specification error". *Econometrica: Journal of the econometric society*, pp. 153–161.

Hemphill, C Scott and Tim Wu (2020). "Nascent competitors". *University of Pennsylvania Law Review*, pp. 1879–1910.

Henderson, Rebecca M and Kim B Clark (1990). "Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms". *Administrative science quarterly*, pp. 9–30.

Higgins, Matthew J and Daniel Rodriguez (2006). "The outsourcing of R&D through acquisitions in the pharmaceutical industry". *Journal of financial economics* 80.2, pp. 351–383.

Holmes, Thomas J, David K Levine, and James A Schmitz Jr (2012). "Monopoly and the incentive to innovate when adoption involves switchover disruptions". *American Economic Journal: Microeconomics* 4.3, pp. 1–33.

Huntington-Klein, Nick (2021). *The effect: An introduction to research design and causality*. Chapman and Hall/CRC.

Hynes, Niki and Andrew David Elwell (2016). “The role of inter-organizational networks in enabling or delaying disruptive innovation: A case study of mVoIP”. *Journal of Business & Industrial Marketing* 31.6, pp. 722–731.

Ivaldi, Marc, Nicolas Petit, and Selcukhan Unekbas (2023). “Killer Acquisitions: Evidence from EC Merger Cases in Digital Industries”.

Jaffe, Adam, Manuel Trajtenberg, and Rebecca Henderson (1993). “Geographic localization of knowledge spillovers as evidenced by patent citations”. *the Quarterly journal of Economics* 108.3, pp. 577–598.

Jaravel, Xavier, Neviana Petkova, and Alex Bell (2018). “Team-specific capital and innovation”. *American Economic Review* 108.4-5, pp. 1034–1073.

Jorgenson, Dale W (2001). “Information technology and the US economy”. *American Economic Review* 91.1, pp. 1–32.

Jovanovic, Boyan and Peter L Rousseau (2005). “General purpose technologies”. *Handbook of economic growth*. Vol. 1. Elsevier, pp. 1181–1224.

Kapoor, Rahul and Kwanghui Lim (2007). “The impact of acquisitions on the productivity of inventors at semiconductor firms: A synthesis of knowledge-based and incentive-based perspectives”. *Academy of Management Journal* 50.5, pp. 1133–1155.

Katz, Michael L (2021). “Big Tech mergers: Innovation, competition for the market, and the acquisition of emerging competitors”. *Information Economics and Policy* 54, p. 100883.

Kim, J Daniel (2020). “Startup acquisitions as a hiring strategy: Worker choice and turnover”. *Available at SSRN 3252784*.

Kim, Jieun, Yongtae Park, and Youngjo Lee (2016). “A visual scanning of potential disruptive signals for technology roadmapping: investigating keyword cluster, intensity, and relationship in futuristic data”. *Technology Analysis & Strategic Management* 28.10, pp. 1225–1246.

King, Gary and Richard Nielsen (2019). "Why propensity scores should not be used for matching". *Political Analysis* 27.4, pp. 435–454.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard (2019). "Children and gender inequality: Evidence from Denmark". *American Economic Journal: Applied Economics* 11.4, pp. 181–209.

Koski, Heli A, Otto Kässi, and Fabian Braesemann (2023). "Killers on the Road of Emerging Start-Ups-Implications for Market Entry and Venture Capital Financing". *Available at SSRN 4390546*.

Kuhn, Jeffrey, Kenneth Younge, and Alan Marco (2020). "Patent citations reexamined". *The RAND Journal of Economics* 51.1, pp. 109–132.

Kwoka, Jr John and Tommaso Valletti (2020). "Scrambled eggs and paralyzed policy: breaking up consummated mergers and dominant firms". *Industrial and Corporate Change*.

— (2021). "Unscrambling the eggs: breaking up consummated mergers and dominant firms". *Industrial and Corporate Change* 30.5, pp. 1286–1306.

Lampe, Ryan (2012). "Strategic citation". *Review of Economics and Statistics* 94.1, pp. 320–333.

Lanjouw, Jean O and Mark Schankerman (2001). "Characteristics of patent litigation: a window on competition". *RAND journal of economics*, pp. 129–151.

— (2004). "Patent quality and research productivity: Measuring innovation with multiple indicators". *The Economic Journal* 114.495, pp. 441–465.

Lemley, Mark A and Matthew Wansley (2024). "Coopting Disruption". *Available at SSRN 4713845*.

Leone, Andrew J, Steve Rock, and Michael Willenborg (2007). "Disclosure of intended use of proceeds and underpricing in initial public offerings". *Journal of Accounting Research* 45.1, pp. 111–153.

Lerner, Joshua (1994). "The importance of patent scope: an empirical analysis". *The RAND Journal of Economics*, pp. 319–333.

Lerner, Joshua, Morten Sorensen, and Per Strömberg (2011). "Private equity and long-run investment: The case of innovation". *The Journal of Finance* 66.2, pp. 445–477.

Lodh, Suman and Maria Rosa Battaggion (2015). "Technological breadth and depth of knowledge in innovation: the role of mergers and acquisitions in biotech". *Industrial and Corporate Change* 24.2, pp. 383–415.

Lugani, V. (2023). *A timeline for startup M&A processes: Key steps and factors to consider*. URL: https://techcrunch.com/2023/01/10/a-timeline-for-startup-ma-processes-key-steps-and-factors-to-consider/ (visited on 11/12/2023).

Lunden, I. (2015). *Google confirms acquisition of Agawi, a specialist in streaming native mobile apps.* URL: https://techcrunch.com/2015/06/18/report-last-year-google-secretly-acquired-agawi-a-specialist-in-streaming-native-mobile-apps/.

Marco, Alan (2007). "The dynamics of patent citations". *Economics Letters* 94.2, pp. 290–296.

Martin, Jonh F. (2015). "The Myth of the 18-Month Delay in Publishing Patent Applications". URL: https://ipwatchdog.com/2015/08/03/the-myth-of-the-18-month-delay-in-publishing-patent-applications/id=60185/.

Melamed, Douglas (May 2021). "Dynamic Antitrust Discussion Series: Killer Acquisitions". Virtual Webinar. URL: https://itif.org/events/2021/05/07/dynamic-antitrust-discussion-series-killer-acquisitions.

Metz, Cade et al. (2023). *Ego, Fear and Money: How the A.I. Fuse was Lit*. Accessed on August 14, 2024. URL: https://www.nytimes.com/2023/12/03/technology/ai-openai-musk-page-alt-man.html.

Michaels, Ed, Helen Handfield-Jones, and Beth Axelrod (2001). *The war for talent*. Harvard Business Press.

Momeni, Abdolreza and Katja Rost (2016). "Identification and monitoring of possible disruptive technologies by patent-development paths and topic modeling". *Technological Forecasting and Social Change* 104, pp. 16–29.

Moraga-González, José Luis, Evgenia Motchenkova, and Saish Nevrekar (2022). "Mergers and innovation portfolios". *The RAND Journal of Economics* 53.4, pp. 641–677.

Morzenti, Giovanni (2022). *Antitrust policy and innovation*. Tech. rep. Mimeo, University of Bocconi.

Motta, Massimo and Martin Peitz (2021). "Big tech mergers". *Information Economics and Policy* 54, p. 100868.

Motta, Massimo and Emanuele Tarantino (2016). "The effect of a merger on investments". *CEPR Discussion Paper No. DP11550*.

Narin, Francis and Elliot Noma (1985). "Is technology becoming science?" *Scientometrics* 7.3, pp. 369–381.

Park, Michael, Erin Leahey, and Russell J Funk (2023). "Papers and patents are becoming less disruptive over time". *Nature* 613.7942, pp. 138–144.

Parker, Geoffrey, Georgios Petropoulos, and Marshall Van Alstyne (2021). "Platform mergers and antitrust". *Industrial and Corporate Change* 30.5, pp. 1307–1336.

Pforr, Klaus (2014). "Femlogit—implementation of the multinomial logit model with fixed effects". *The Stata Journal* 14.4, pp. 847–862.

Pierce, P. (2022). *How Long Does it Really Take to Sell a Business?* URL: https://www.madisonstreetcapital.com/how-long-does-it-really-take-to-sell-a-business/ (visited on 11/14/2023).

Polsky, Gregg D. and John F. Coyle (2013). "Acqui-hiring." *Duke Law Journal* 63.2, p. 281.

Prado, Tiago S and Johannes M Bauer (2022). "Big Tech platform acquisitions of start-ups and venture capital funding for innovation". *Information Economics and Policy* 59, p. 100973.

Puranam, Phanish, Harbir Singh, and Maurizio Zollo (2006). "Organizing for innovation: Managing the coordination-autonomy dilemma in technology acquisitions". *Academy of Management journal* 49.2, pp. 263–280.

Ranft, Annette L and Michael D Lord (2000). "Acquiring New Knowledge: the Role of Retaining Human Capital".

— (2003). "Acquiring new technologies and capabilities: A grounded model of acquisition implementation". *Quality control and applied statistics* 48.5, pp. 581–584.

Régibeau, Pierre and Katharine Rockett (2010). "Innovation cycles and learning at the patent office: does the early patent get the delay?" *The Journal of Industrial Economics* 58.2, pp. 222–246.

Reinganum, Jennifer F (1983). "Uncertain innovation and the persistence of monopoly". *The American Economic Review* 73.4, pp. 741–748.

Righi, Cesare and Timothy Simcoe (2019). "Patent examiner specialization". *Research Policy* 48.1, pp. 137–148.

Röller, Lars-Hendrik, Johan Stennek, and Frank Verboven (2006). "Efficiency gains from mergers". *European merger control: do we need an efficiency defence*, pp. 84–201.

Rosenkopf, Lori and Paul Almeida (2003). "Overcoming local search through alliances and mobility". *Management science* 49.6, pp. 751–766.

Rosenkopf, Lori and Atul Nerkar (2001). "Beyond local search: boundary-spanning, exploration, and impact in the optical disk industry". *Strategic management journal* 22.4, pp. 287–306.

Rysman, Marc and Timothy Simcoe (2008). "Patents and the performance of voluntary standard-setting organizations". *Management science* 54.11, pp. 1920–1934.

Sampat, Bhaven N (2010). "When do applicants search for prior art?" *The Journal of Law and Economics* 53.2, pp. 399–416.

Scherer, Frederic M and Dietmar Harhoff (2000). "Technology policy for a world of skew-distributed outcomes". *Research policy* 29.4-5, pp. 559–566.

Schimmer, M, G Mueller-Stewens, and P Sponland (2010). "The battle between Apple, Microsoft and Google: Strategic lessons from a converging internet industry 2000–2010". *University of St. Gallen, Case Study*, pp. 310–245.

Schmidt, Hedvig (2023). "Merger regulation in the digital economy and the forgotten goal of innovation". *Journal of Antitrust Enforcement* 11.Supplement_1, pp. i102–i126.

Scott Morton, Fiona et al. (2019). "Committee for the study of digital platforms: Market structure and antitrust subcommittee report". *Chicago: Stigler Center for the Study of the Economy and the State* 36.

Seitz, Nikolaus and Jasmin Kopp (2023). "Does Acqui-Hiring Really Pay Offs? An Empirical Investigation of Founders' Retention". *Academy of Management Proceedings*. Vol. 2023. 1. Academy of Management Briarcliff Manor, NY 10510, p. 15156.

Selby, Jaclyn and Kyle J. Mayer (2013). "Startup firm acquisitions as a human resource strategy for innovation: The acqhire phenomenon." *Academy of Management Proceedings* 2013.1, p. 17109.

Shane, Scott (2001). "Technological opportunities and new firm creation". *Management science* 47.2, pp. 205–220.

Shang, Shengwu, Erik Nesson, and Maoyong Fan (2018). "Interaction terms in poisson and log linear regression models". *Bulletin of Economic Research* 70.1, E89–E96.

Shapiro, Carl (2011). "Competition and innovation: did arrow hit the bull's eye?" *The rate and direction of inventive activity revisited*. University of Chicago Press, pp. 361–404.

Shelegia, Sandro and Massimo Motta (2021). "The" kill zone": Copying, acquisition and start-ups' direction of innovation".

Si, Steven and Hui Chen (2020). "A literature review of disruptive innovation: What it is, how it works and where it goes". *Journal of Engineering and Technology Management* 56, p. 101568.

Simon, Phil and M Joel (2011). "The age of the platform: How Amazon, Apple, Facebook, and Google Have Redefined Business". *Las Vegas: Motion Publishing*.

Squicciarini, Mariagrazia, Hélène Dernis, and Chiara Criscuolo (2013). "Measuring patent quality: Indicators of technological and economic value". *OECD Science, Technology and Industry Working Papers* 3.

Statista (2024). *Most popular multi-platform web properties in the U.S. based on number of unique visitors.* URL: `https://www.statista.com/statistics/271412/most-visited-us-web-properties-based-on-number-of-visitors/`.

"Stigler Report" (2019). *Stigler Committee on Digital Platforms, Final Report*.

Stock, James H, Jonathan H Wright, and Motohiro Yogo (2002). "A survey of weak instruments and weak identification in generalized method of moments". *Journal of Business & Economic Statistics* 20.4, pp. 518–529.

Sun, Liyang and Sarah Abraham (2021). "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects". *Journal of econometrics* 225.2, pp. 175–199.

Sutton, John (2001). *Technology and market structure: theory and history*. MIT press.

Tegernsee (2012). "Study Mandated by the Tegernsee Heads: 18-Month Publication". URL: `https://www.uspto.gov/sites/default/files/ip/global/18_months_publication.pdf`.

Thompson, Peter and Melanie Fox-Kean (2005). "Patent citations and the geography of knowledge spillovers: A reassessment". *American Economic Review* 95.1, pp. 450–460.

"TomTom/Tele Atlas" (2008). *Case M.4854, Commission decision of 14 May 2008, paragraphs 244-250*.

Tong, Xuesong and J Davidson Frame (1994). "Measuring national technological performance with patent claims data". *Research Policy* 23.2, pp. 133–141.

Trajtenberg, Manuel (1990). "A penny for your quotes: patent citations and the value of innovations". *The Rand journal of economics*, pp. 172–187.

Trajtenberg, Manuel, Rebecca Henderson, and Adam Jaffe (1997). "University versus corporate patents: A window on the basicness of invention". *Economics of Innovation and new technology* 5.1, pp. 19–50.

Tushman, Michael L and Philip Anderson (2018). "Technological discontinuities and organizational environments". *Organizational innovation*. Routledge, pp. 345–372.

Valletti, Tommaso and Hans Zenger (2019). "Increasing market power and merger control". *CLPD* 5, p. 40.

Varian, Hal R (2021). "Seven deadly sins of tech?" *Information Economics and Policy* 54, p. 100893.

Verginer, Luca et al. (2022). "The Impact of Acquisitions on Inventors' Turnover in the Biotechnology Industry". *arXiv preprint arXiv:2203.12968*.

WIPO (2022). "PCT Yearly Review 2022". URL: https://www.wipo.int/edocs/pubdocs/en/wipo-pub-901-2022-en-patent-cooperation-treaty-yearly-review-2022.pdf.

Wooldridge, Jeffrey M (2010). *Econometric analysis of cross section and panel data*. MIT press.

— (2023). "Simple approaches to nonlinear difference-in-differences with panel data". *The Econometrics Journal* 26.3, pp. C31–C66.

Xiao, Jing and Åsa Lindholm Dahlstrand (2023). "Skill-biased acquisitions? Human capital and employee mobility in small technology firms". *Small Business Economics* 60.3, pp. 1219–1247.

Zheng, Lu, Klaus Ulrich, and Javier Sendra-García (2021). “Qualitative comparative analysis: Configurational paths to innovation performance”. *Journal of business research* 128, pp. 83–93.

Zucker, Lynne G, Michael R Darby, and Marilynn B Brewer (1998). “Intellectual human capital and the birth of US biotechnology enterprises. American Economics Review 88 (1) 290–306”. *Intellectual Capital and Financing Decisions: Evidence from the US Patent Data* 23.8.