

Big Tech Acquisitions and Innovation: An Empirical Assessment ¹

Laureen de Barsy^{2*} and Axel Gautier^{3*}

*HEC Liège, LCII, Université de Liège

January 27, 2024

PRELIMINARY VERSION – Comments welcome

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 on their evolution after acquisition. This paper intends to fill in this gap by looking at the development of the technologies owned by the acquired firms. We focus on technologies protected by a patent and we investigate whether an acquisition by a big tech contributes to their development. For this analysis, we use patent citations as a proxy for the innovation effort by the acquirer. Our main result is to show that acquisition increases the innovation effort of the acquirer but only temporarily. After 1.5 year, there is no longer a significant impact of the acquisition on the acquirer's innovation effort. This decline is relatively larger when the acquired patent belongs to a core technology field of the acquiring firm or to a large patent portfolio. On the contrary, citations by the rest of the industry are not negatively affected by acquisition, which does not corroborate the idea that the acquired technology has reached its maturity.

¹We are grateful to Laurie Ciaramella, Sylvain Dejean, Xavier D'Haultfoeuille, Songyuan Ding, Luise Eisfeld, Roxana Fernandez, Malka Guillot, Giovanni Morzenti, Valerio Sterzi, Maren Ulm and Jeffrey Wooldridge for their helpful comments, as well as to the ARC project.

²ldebarsy@uliege.be

³agautier@uliege.be

I Introduction

One of the most notable transformations of our economy over the last 30 years is its move towards digitalization. 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 2000s, these companies are now 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 considered¹, especially given the very high rate at which these platforms acquire start-ups. In an interview on CNBC, Tim Cook (2019), 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."* Despite this intense rate of acquisition, very few have been reviewed by a competition authority² and, up to date, only one of them has ever been blocked (CMA 2021). This can first be explained by the fact that most of these transactions do not meet the turnover-based notification thresholds to be subject to a review by a competition authority. Second, competition authorities are in charge of controlling a market that is becoming more complex and opaque every day, and over which 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).

It is important that competition authorities better scrutinize the acquisitions by the GAFAM in order to properly assess their impacts on competition. Proposals are being made to reinforce the control of mergers by competition authorities, especially in the digital sector.³

¹See for instance Argentesi et al. 2021; Crémer, Montjoye, and Schweitzer 2019; "Stigler Report" 2019.

²More than 97% of acquisitions in the technology sector have reportedly never been subject to scrutiny by a competition authority (Kwoka and Valletti 2020).

³Reforms to the merger control framework in the digital economy are being discussed throughout the world. In response to the low number of transactions actually examined by a competition authority, some countries have already introduced a reform of the legal notification thresholds. For example, Austria and

In a static environment absent efficiencies and synergies, economic theory predicts that, by relaxing a competitive pressure, horizontal mergers necessarily lead 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 GAFAM are spending billions in R&D and many of the firms they acquired are young and innovative startups that often develop new technologies.

A merger can have a positive effect on innovation, and this could be used as an argument in the “balance of harm” approach of competition authorities. In practice, 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, the European Commission recognized that the merger between a navigation systems provider and a digital maps developer would allow to deliver “better maps - faster” (TomTom/Tele Atlas 2008). These efficiencies would thus translate into the acquired technology being further developed after acquisition.

However, these positive benefits of the merger on innovation are far from being granted⁴ and a merger may have a negative impact on innovation. In some case, mergers are 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 acquisitions of a startup by a dominant platform results in a strengthening of its dominance, a reduction of effective competition, and a loss of innovation (Motta and Peitz 2021). These concerns are even growing given the high number of acquisitions and the lack of information on the evolution of the startups after they have been acquired.

The existing evidence show 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

Germany now apply a notification threshold based on the transaction price. Since March 2021, the European Commission has also allowed Member States to refer to it the examination of transactions that do not meet the turnover thresholds when the latter does not reflect the actual or future competitive potential of at least one of the merging parties (European Merger Regulation, Art. 22). Some experts also envisage a “reversal of the burden of proof”, whereby the merging parties would bear the burden of demonstrating the absence of anti-competitive effects (“Stigler Report” 2019).

⁴In the Dow/DuPont merger case, the European Commission expressed concerns that the merger would have reduced innovation. The parties agreed to divest assets in overlapping markets to preserve the industry’s incentives to innovate.

acquisition. Affeldt and Kesler 2021 focus on merging involving ‘apps’ and they document that half of those apps were discontinued after an acquisition by a tech giant. Einfeld 2022 studies startup acquisition in the software industry and finds that 57% of the acquired products have been discontinued under their original brand name after acquisition. Product discontinuation is particularly a concern when the target is small (Gautier and Lamesch 2021).⁵ However, a project 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 and this paper intends to fill in 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, materialized 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 patent protecting that technology. This implies that the development of a technology is materialized 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 after acquisition translates into more (less) citations to the acquired patents. This is precisely, the relation with intend to test.

For our analysis, we construct a sample of firms acquired by one of the GAFAM since 1996 and we identified those that have been granted 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

⁵Ivaldi, Petit, and Unekbas 2023 do not identify a product discontinuation in their review of the (larger-scale) digital mergers investigated by the European Commission.

derive from a low probability of imitation, high costs of patenting (e.g. administrative costs and renewal fees), length of the grant procedure⁶, 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 68% of the 96 biggest acquired firms (i.e. with a total funding above \$10 million)⁷.

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 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, by means of a propensity score weighting design, we construct a set of patents with comparable characteristics but that have not been acquired. We can 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. This allows us to identify the dynamic effects of acquisition on the development of acquired technologies.

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.

Next, we test whether this observed pattern is identical for all technologies and all patents. First, we find that the boost in citations is higher for technologies that are more novel, which is a rather intuitive result but, even for these novel technologies, the effect seems to fade away in the longer term. Second, the effect of acquisition on inno-

⁶US patents take approximatively 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).

⁷Based on funding data retrieved from Crunchbase.

vation is stronger and more persistent for patents belonging to relatively small patent portfolios. For patents belonging to a large portfolio, we do not observe a significant boost in citations but the slowing down after 1.5 year does remain. This result could be explained by different drivers for the acquisition; while the acquisition of firms with small patent portfolios are likely technology-driven, larger portfolios would rather be acquired for their other assets, like products, clients, network or talent. Third, we distinguish between acquired technologies belonging to a field in which the acquirer holds a strong position (what we call “core” technologies), and technologies outside of the acquirer’s core fields (“peripheral” technologies). Our results show that the boost in citations is mainly driven by peripheral technologies, while the slowing down in the acquired technology developments appears to be driven by core technologies. This suggests a potential competitive explanation for our results. The acquisition of peripheral technologies could be motivated by their R&D potential. Outside of their main technology fields, big techs seem to acquire technologies, engineers and labs, while acquisitions in their main technology fields may have a more strategic motive. In those core fields, Big Tech seems to acquire firms that could represent a competitive threat to its own research and shelve these technologies after they have been acquired. Such a different strategy may explain the difference in citations trends for core and peripheral technologies.

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 being commercialized. 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 is unlikely to explain Big Tech’s declining interest for the development of acquired technologies.

Related literature

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

the acquired company. The earlier literature focused on the intensity of the innovation effort but, more recently, the literature also focuses on the direction of innovation.

The start-up innovative effort can first be impacted through the possibility of buy-out. 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. To do so, the start-up would distort its portfolio of projects towards the interests of the platform such as to maximize the probability of being acquired and the expected acquisition rents (Bryan and Hovenkamp 2020b, Dijk, Moraga-González, and Motchenkova 2021). This leads to less radical innovation and lower quality (Cabral 2018, Katz 2021) but it may also stimulate the innovation effort (Motta and Peitz 2021). Furthermore, digital platforms may engage in exclusionary practices, for instance by reducing interoperability with the startup's product or by imitating its main features and this threat will drive startups away from the platforms' core market (Motta and Peitz 2021, Shelegia and Motta 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&A 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&A 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 internalize 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 internalize these externalities (Federico, Langus, and Valletti 2018). 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).⁸

The main driver of the negative effects of M&A on innovation is their impact on the market structure. According to 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 internalize 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&A 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&A 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).

A start-up might also not have the resources to bring the project to the market and the acquisition by a large platform may bring the necessary resources to complete the

⁸If 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.

project. However, the acquirer 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” the innovation.

Several papers have tried to assess empirically the impact of mergers on innovations, 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 merger in digital industries. Doan and Mariuzzo 2022 analyze 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 analyze 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 analyze the impact of the merger at technology/patent level, as we do, but at a more aggregated 'market' level. For that they 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 found a significant negative impact of acquisitions on venture-capital funding. The effect on innovation is less clear cut and

it depends on the period and on the acquirer. The initial negative effect observed for mergers before 2010 becomes positive for mergers after this date, with a different magnitude for each GAFAM, the effects, both positive and negative, being the strongest 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 big tech in a given industry increases the venture capital activity in that industry with a significant increase in the number of deals and funding. However, they 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 2, we describe the main features of Big Tech acquired technologies (2.1) and the construction of our working datasets (2.2). Section 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 4 and our main results in Sections 5.1 and 5.2, with tests of robustness in Section 5.3. We develop additional analyses and extensions in Sections 6 and 7, and Section 8 concludes.

2 Empirical Methodology

For our analysis, we construct two samples of patents. The first is a sample 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. In this section, we describe the data collection and the construction of the working datasets.

2.1 Data and Variables

2.1.1 Big Tech acquisitions

Our working sample is constructed in three steps, as presented in Table 1.

We first create a dataset of firm's acquisitions by Alphabet, Amazon, Apple, Meta and Microsoft. To obtain as complete of a list of Big Tech acquisitions as possible, we merge 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).⁹ We retrieve information on the identities of the acquired firms and on the dates at which their acquisitions were announced. On this basis, we identify 859 public big tech acquisitions closed between January 1996 and January 2021 (see first column of Table 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. We focus on patent filed before acquisition¹¹ that have later been granted. Because there is a lag between the filing date and the granting date, an acquired patent could be granted

⁹We 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 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.

¹⁰USPTO-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). 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.

¹¹Patents filed under the target's name after acquisition are considered as filed by the acquirer.

after acquisition.¹² We find that 273 of these firms have filed at least one patent application, of which 252 before being acquired (see second and third columns of Table 1).

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 rises to 76% when we consider the biggest firms (i.e. with a total funding above \$10 million¹³).

Table 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	106	34	27
Apple	128	53	52
Facebook	104	18	18
Google	264	75	67
Microsoft	257	93	88
TOTAL	859	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.

2.1.2 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 retrieve information on the patent age based on its filing date.

Core technology Another interesting information included in the Patent Views database is the technology fields to which a patent belongs. This information is recorded

¹²Before 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).

¹³Based on funding data retrieved from Crunchbase and Orbis.

in the CPC classification, which contains 131 subsections at the two-digits level (detailed list). On this basis, we will be able to explore the potential relation between the technological content of an acquirer’s portfolio and its target’s portfolio.

We first compute the frequency of each of the 131 CPC subsections for all US patents at the yearly level. Next, we compute the share represented by each tech giant’s portfolio in these respective subsections. We consider that a given CPC subsection represents a Big Tech core technology field, in a given year, if the Big Tech’s portfolio contains at least 1% of all occurrences of that technology in that year. Finally, we identify acquired patents associated with at least one of their acquirer’s core technology in the year of their acquisition. These patents are marked as “core”.

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 by the acquiring firm to the acquired technology reflect whether the technology is being further improved by its acquirer. Appendix I discusses the potential limitations attached to this use of patent citations data.

We obtain information on forward citations by the acquiring firm by taking the following steps. First, and in addition to the sample of Big Tech-acquired patents described in the previous section, we identify all granted patents filed by Big Tech itself. Next, we retrieve the application identifiers of all patents containing a citation to a patent filed by a Big Tech-acquired firm from the Patent Views database.¹⁴ Patents cited by their acquirer can then be identified by matching these application identifiers of the citing patents to the filing firms. 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.¹⁵

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’. Citations data is available until July 2022; 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 2017, 5 years before the data collection. Figure 8 in Appendix II shows that, from 2018

¹⁴The Patent Views database covers all citations made by US granted patents.

¹⁵We assume citations are observed from the date of filing.

onwards, citing patents start being less likely to appear in the Patent Views database because they have not been granted yet. Our choice of ending the study period in June 2017 is therefore conservative.

2.1.3 Acquired firms

In the end, for each patent in our database, we can identify the acquirer, the timing of acquisition, the patent's age, the number of forward citations made every month to this patent, and whether it belongs to a technology field in the acquirer's core business at acquisition. We construct a dataset containing all patents belonging to one of the 252 Big Tech-acquired firms, and we select those firms that have published, pre-acquisition, at least one patent further cited by their acquirer in a patent filed before July 2017. We end up with a working sample of 143 firms, i.e. 143 patent portfolios. Table 2 presents summary statistics of these data samples.

Table 2: Big Tech acquired patents portfolios

	Firms Count	Portfolio size (patents #)		Patent age at acquisition (y)		Aquirer core technology
		Mean	SD	Mean	SD	% Patents
Big Tech acquired portfolios						
AMZN	27	22.07	64.62	3.31	2.59	65%
APPL	52	14.21	19.72	4.00	2.59	39%
FCBK	18	7.56	17.68	4.31	4.37	81%
GOOG	67	30.98	143.16	3.90	2.14	24%
MSFT	88	16.52	51.88	3.86	2.79	80%
TOTAL	252	19.84	83.12	3.87	2.71	49%
Big Tech acquired portfolios cited by their acquirer before July 2017						
AMZN	12	15.25	24.81	3.00	2.45	10%
APPL	29	19.79	23.49	4.83	3.22	38%
FCBK	6	5.17	4.88	3.16	3.28	47%
GOOG	35	56.66	195.84	5.04	2.38	27%
MSFT	61	22.57	61.46	4.52	3.58	88%
TOTAL	143	29.01	105.83	4.52	3.17	56%

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

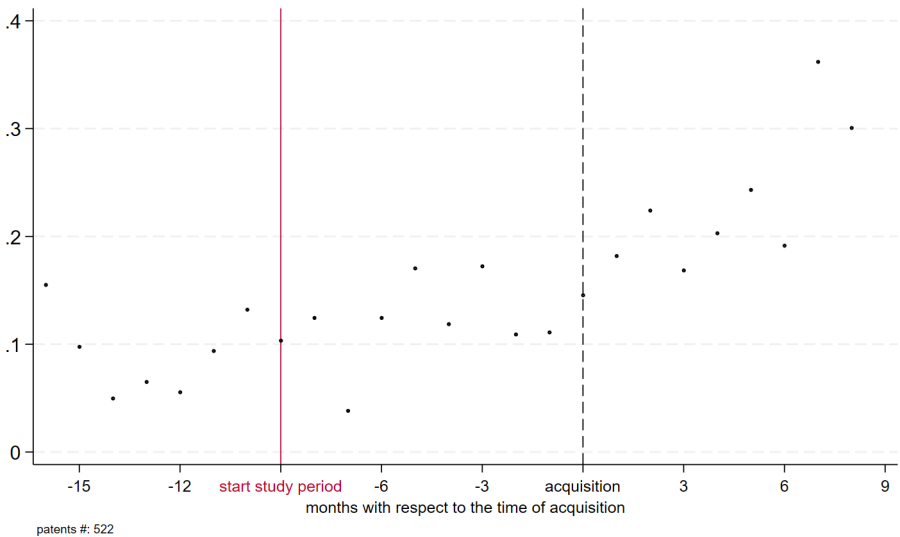
2.2 Working sample

2.2.1 Acquired patents

The next step is to construct a balanced panel of observations and we select patents that we observe every month during 4 years around acquisition. In our model, the event time is 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,¹⁶ 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. To meet these two goals, we choose a period of 9 months before the acquisition announcement (see Figure 1), which allows to observe the evolution of citations before acquisition while including targets of all ages.

Figure 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.

¹⁶Firms 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).

For the post-treatment, we choose a period of 3 years after the acquisition announcement. 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 859 Big Tech acquisitions.

Of all acquired patents observed in this 4 years-window, 541 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.

2.2.2 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.¹⁷ 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.

¹⁷A similar study design is developed by Kleven, Landais, and Søggaard (2019), who assign placebo births to individuals who never had children by drawing from the observed distribution of age at first child among parents.

The first column of Table 3 presents the number of patents in these two groups: 541 patents (accounting for 80 firms) undergo an acquisition event, and 70,136 are assigned a placebo acquisition date. The next columns of Table 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 8.63 citations/acquired patent against 3.53/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 V).

Table 3: Observations over the whole study period

	Cited patents	Patent 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 underrepresented. 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-treatment, 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 III.

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 exogeneity of the acquisition event, as well as 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.

3.1 Baseline: Sharp event study

For our baseline model, we adopt a sharp event study approach as developed by Kleven, Landais, and Søggaard (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 $t = 0$.¹⁸ 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 (i.e. the number of forward citations might depend on the stage of a patent’s life)

¹⁸The 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: $t \in \{-3 = (-10m, -9m, -8m), \dots, 0 = (-1m, 0m, 1m), \dots\}$ with $0m$ when the filing month coincides with acquisition.

and business-cycle trends (i.e. 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 firm j at event time t and date d . Target-specific fixed effects are captured by $firm_j$. We control non-parametrically 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' ($d = 1996q1, 1996q2, \dots, 2017q2$).¹⁹ 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:

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

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 $f(\cdot)$, 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 (see Appendix V). 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). 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 is significantly different from zero ($\chi^2 = 2985$), which contradicts the assumption of the Poisson model. On this basis, we can confirm that a Negative Binomial regression should be used.

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

¹⁹The calendar date dummy is defined as the quarter associated with the month in which the citing patent is filed, e.g. $(2013m7, 2013m8, 2013m9) = 2013q3$.

as:

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

where $\mu(\cdot)$ 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 $Cit_{p,j,t,d} = 0, 1, 2, \dots$

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, $\hat{\theta}_t^1$ identifies the expected difference in log citations between quarter t and the reference group ($t = 0$): $\hat{\theta}_t^1 = \ln(Cit_{p,j,t,d} \mid age_{p,d}, d, firm_j) - \ln(Cit_{p,j,0,d} \mid age_{p,d}, d, firm_j)$. To obtain a more intuitive interpretation of our results, we will use the incident rate ratios: $e^{\hat{\theta}_t^1} = \frac{Cit_{p,j,t,d} \mid age_{p,d}, d, firm_j}{Cit_{p,j,0,d} \mid age_{p,d}, d, firm_j}$. 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 IV.

3.2 Difference-in-semielasticities

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 standard normal distribution (as described in Section 2.2.2). We rewrite model 1 as follows:

$$Cit_{p,t,d} = f(J'\theta^2, A_p t^1, J' A_p \alpha^1, age_{p,d} \beta^2, M' \gamma^2) \quad (2)$$

where $A_p = 1$ if patent p is acquired, $A_p = 0$ otherwise.²⁰

²⁰In this second model specification, firms fixed effects are no longer accounted for as we cannot retrieve the identities of all firms cited by Big Tech patents.

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 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),²¹ i.e. the acquisition status’ impact on the semielasticity of citations with respect to the event time: $e^{(\widehat{\theta}_i^+ + \widehat{\alpha}_i^+)} - e^{(\widehat{\theta}_i^-)}$. The validity of the identification parallel trends assumption can be verified from the pre-acquisition DIS.

4 Preliminary analysis

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

If we look at the describe statistics, we observe that an acquired 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 control for life-cycle and business-cycle trends, we define simplified version of models 1 and 2, with the dummy variable *Post* taking the value 1 after acquisition:

$$Cit_{p,j,d} = f(Post \theta^1, firm_j \xi^1, age_{p,d} \beta^1, M' \gamma^1) \quad (3)$$

$$Cit_{p,d} = f(Post \theta^2, A_p \iota^1, Post A_p \alpha^1, age_{p,d} \beta^2, M' \gamma^2) \quad (4)$$

The estimation results are presented in Table 4. Column 1 presents the estimation based on the sample of acquired patents only (Model 3). The results show a significant

²¹When 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).

increase in citations after acquisition. The model estimates that an acquired patent receives 35% ($IRR = e^{\hat{\theta}^1} = 1.35^{***}$ (0.10)) more citations by its acquirer after the acquisition. The results for Model 4 on (unweighted) acquired and non-acquired patents are similar, with an estimated citations increase of 48% ($DIS = e^{(\hat{\theta}^2 + \hat{\alpha}^1)} - e^{\hat{\theta}^2} = 0.48^{***}$ (0.12)) after acquisition. 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.

Table 4: Big Tech citations to acquired patents

	Model (3)	Model (4)
Post	.30 ^{***} (.07)	.15 ^{***} (.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

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

5 Impact of acquisition on the acquired technology

We present below estimates of the impact of a big tech acquisition on the development of the acquired technology as measured by citations to the associated patents. Model 1 is estimated on the balanced panel of Big Tech-acquired patents. Model 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.

5.1 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 are presented on Figure 2. On the figure, we represent the estimated incident rate ratios ($e^{\hat{\theta}_t}$) 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 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 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. 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.

5.2 Difference-in-semielasticities

We present on Figure 3 the DIS estimated based on Model 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 non-parametrically for life-cycle and business-cycle trends, for acquired patents with respect to non-acquired patents. A value above 0 means that citations of acquired patents are higher relative to non-acquired ones.

Figure 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: $e^{(\hat{\theta}_t^2 + \hat{\alpha}_t^1)} - e^{(\hat{\theta}_t^2)}$ from model 2. These coefficients are estimated on a balanced sample of patents in a 4 year-window around (simulated) acquisition.

In support of the assumption that citations for 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.

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

many alternative specifications.

5.3 Robustness checks

To test for the robustness of our results, we replicate our baseline analysis with more citations determinants included as regressors in the model and with alternative study periods. These robustness checks are presented in Appendix VI.

First, we propose to replicate our analysis based on alternative specifications including more citations determinants. 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.

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 and last, we include *Motorola Mobility*, which was acquired by Google but later sold to Lenovo and hence not included in our baseline sample.

In all our specifications, we found results that are consistent with those presented above, with an initial boost in citations, followed by a slowdown.

6 Extensions

In this section, we extend our baseline results and we test our models with different subsamples. We include some target's characteristics that we believe can influence patent citations. First, we control for the technical novelty of the acquired patents, second, we control for the portfolio size of the acquired firm and, last, we control for the proximity of the acquired patents with the acquirer's patent portfolio.

6.1 Technical novelty

Big tech may have different strategies for different patent types. In particular, they may selectively develop acquired patents and put more innovation effort in patents that have more potential and which are path-breaking innovation. For that, we use a measure of patent novelty developed by Arts, Hou, and Gomez 2021. They identify pairwise

combinations of keywords in the title, abstract, or claims of a patent introduced for the first time in history by granted U.S. utility patents. They find that these new combinations of keywords outperform the traditional novelty measures based on patent classification and citations to measure technical novelty at the time of filing.²²

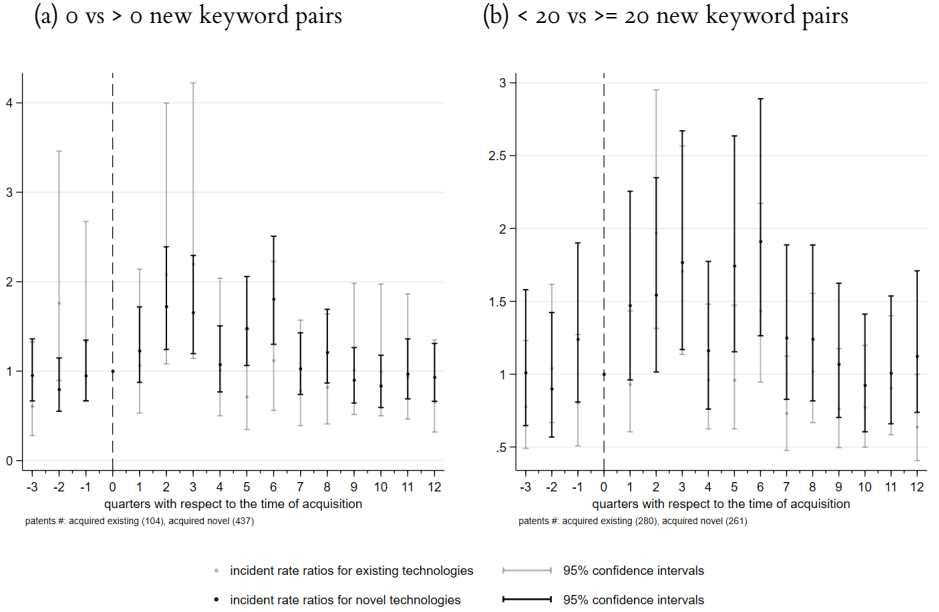
In our sample, 80% of the acquired patents have at least one new keywords pair and 46% have 20 or more new keywords pairs. We define a dummy variable based on these two measures of novelty: $Novel_p$ takes the value 1 if the patent p contains at least one (/at least 20) new keyword pair(s), 0 otherwise. On this basis, we can rewrite model 1 to allow the event time impact to vary with the novelty of the acquired technology:

$$Cit_{p,j,t,d} = f(J'\theta^4, Novel_p\zeta^1, J'Novel_p\eta^1, age_{p,d}\beta^4, M'\gamma^4, firm_j\xi^2) \quad (5)$$

The results are presented on Figure 4. The figure shows a similar inverse U-shaped trend in citations post-acquisitions. We find that both patent groups, more or less novel, exhibit a boost in citations by their acquirer just after acquisition, followed by a decline after 1.5 year.

²²To determine 'novelty', Arts, Hou, and Gomez 2021 gathered patents associated with prestigious awards like the Nobel Prize and the National Inventor Hall of Fame. These patents are believed to protect highly innovative technologies that have had a significant influence on subsequent technical advancements. 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.

Figure 4: Big Tech citations to acquired patents relative to acquisition, by target's technical novelty



Notes: The graphs show the incident rate ratios from model 5 for existing (e^{θ^4}) and novel ($e^{\theta^4 + \eta^1}$) acquired technologies.

6.2 Portfolio size

As a development to our main results, we want to test whether the effect of acquisition varies with the size of the acquired patents portfolio. To do that, we refine model 1 by allowing the event time impact to vary with the size of the target's patents portfolio:

$$Cit_{p,j,t,d} = f(J'\theta^4, Large_p \zeta^1, J' Large_p \eta^1, age_{p,d} \beta^4, M'\gamma^4, firm_j \xi^2) \quad (6)$$

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

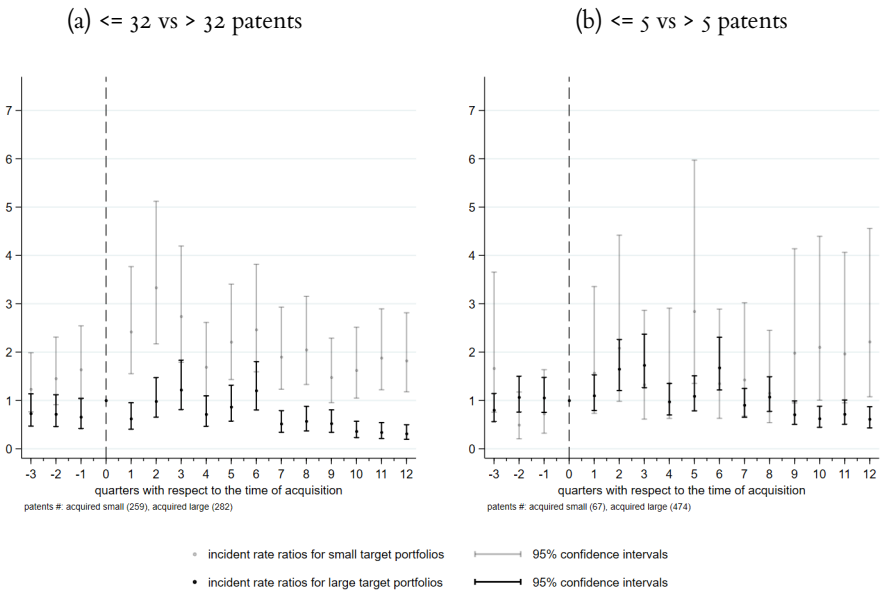
than 5 patents, with a majority of patents falling in this category.

The estimated incident rate ratios are presented on Figure 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 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 for the acquirer.²³ This suggests that the acquisition of small portfolios are more likely to have been driven by a specific patent for which the acquirer exerts a significant effort to further develop it.

²³Let us however remind the reader that patents should be cited at least once by the acquirer to be included in our sample.

Figure 5: Big Tech citations to acquired patents relative to acquisition, by target's size



Notes: The graphs show the incident rate ratios from model 6 for small (e^{θ^1}) and large ($e^{\theta^1 + \eta^1}$) acquired portfolios.

6.3 Core and peripheral technology fields

The results presented above could be interpreted under the prism of the ‘buy vs build’ dilemma. Because there is a likely time lag between the moment a company starts working on a research project and when it files the related patent, the boom in citations after acquisition might relate to some research that had been undertaken before acquisition. At the time, the start-up’s innovative project might have been seen by the platform as a competitive threat. To defend its market, the platform would have invested in developing a substitute project (and thus used the patents protecting the start-up’s technology, which they will need to cite once they file their own patent). If the tech giant fails to replicate the start-up’s technology, it might choose to buy it instead. And if it successfully developed the technology, it might also choose to buy it to eliminate a competitive threat. In both cases, this technology no longer represents a competitive threat since the platform now has a monopoly over it, which would explain the slowing down of the acquired technology development by its acquirer.


To explore this potential interpretations of our results, we want to test whether the effect of acquisition differs depending on the technology to which the acquired patent belongs. The tech giants are acquiring many patents in technology fields where they already hold a strong market position. We identified these patents by defining core technology fields for the acquirer. If Big Tech abandons target's innovative projects because they no longer represent a competitive threat to its own technologies, we should observe that technologies overlapping with their acquirer's core business (i.e. in which the acquirer has focused its own innovative effort) are less likely to be further developed after acquisition.

We allow the event time impact to vary with the technology field to which the acquired technology belongs:

$$Cit_{p,j,t,d} = f(J'\theta^5, Core_p\zeta^2, J'Core_p\eta^2, age_{p,d}\beta^5, M'\gamma^5, firm_j\xi^3) \quad (7)$$

where $Core_p$ takes the value 1 if patent p belongs to a technology field in the acquirer's core business at the time of acquisition, 0 otherwise.

Figure 6: Big Tech citations to acquired patents relative to acquisition, by technology type (acquirer’s core vs peripheral)



Notes: The graph shows the incident rate ratios from model 7 for peripheral (e^{θ^5}) and core ($e^{\theta^5+\eta^2}$) acquired technologies.

Figure 6 presents the coefficient estimates from model 7. When considering each quarter separately, acquirer’s citations to the two technology types (i.e. overlapping with their acquirer’s core business or not) do not seem significantly different. However, when looking at their evolution over event time, this exercise also reveals that the boost in citations just after acquisition is mainly driven by peripheral technologies, i.e. by technologies that do not belong to one of their acquirer’s core businesses. This result suggests that the tech giants use acquisitions to expand to new technological areas, rather than to develop technologies in which they are already strong. In addition, the slowing down in citations after 1.5 year after acquisition is, instead, mainly observed for core technologies. Since we defined a core technology as a patents field of which an acquirer owns at least 1%, we expect this acquirer to face less competition in core technologies than in peripheral technologies. Therefore, this last finding is in line with the ‘buy vs build’ interpretation of our results, which originates in the threat that the target technology represents to its acquirer. Peripheral technologies are unlikely to represent such a threat, since they relate to fields in which the acquirer is less active. Instead, the

slowing down in the development of core technologies after their acquisition could be the sign that their acquirer recovered a comfortable market position.

7 Technology development by the non-acquiring firms

The slow down in citations can be explained by diminishing returns to the innovative effort because the technology reaches its maturity, which would explain that it is subsequently less developed. In 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.

To test for this hypothesis, we look at the evolution of the citations to Big Tech-acquired patents by the other firms in the industry. We can use citations of the acquired patents as a proxy for the research effort to develop the technology by the rest of the industry. According to the technology maturity hypothesis, we should observe a similar slowing down of its development by the industry as a whole. Citations by the acquirer and citations by the rest of the industry should follow a similar pattern.

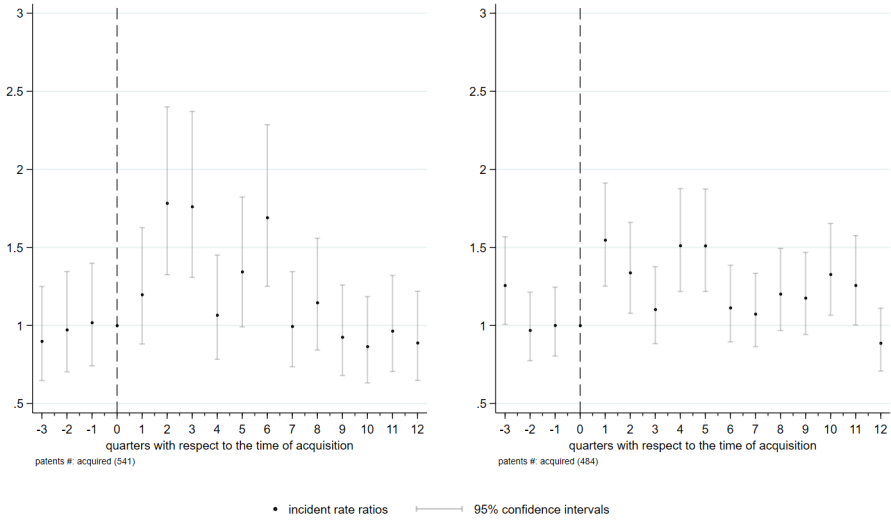
We estimate model 1 on two separate samples: Big Tech-acquired patents cited by their acquirer, and Big Tech-acquired patents cited by other firms than their acquirer. Out of the 541 Big Tech-acquired patents in our sample, 484 are also cited at least once over our study period by other firms than their acquirer. The estimated incident rate ratios ($e^{\hat{\theta}_i^1}$) are presented on Figure 7, separately for these two citing groups.

Figure 7 (a) is the classical inverse U-shaped curve for the citations of the acquirer and Figure 7 (b) represents the citations by the rest of the industry. On Figure 7 (b), we observe that the acquisition induces a positive effect on the citations by the non-acquiring firms; they increase by up to 50% after the acquisition by a tech giant. 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 show a slowdown after 1.5 year, citations by other firms than the acquirer keep increasing up to 2.5 years after acquisition.

Figure 7: 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: $e^{\hat{\theta}_t^1}$ from model 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 technology, eventually more than before its acquisition. This suggests that the improvement potential of the technology has not been exhausted, so technology maturity alone does not seem to provide a credible explanation for the slowing down of the acquired technology development by its acquirer.

Acquisition has a positive impact on the non-merging parties, a result that is consistent with the model of Federico, Langus, and Valletti 2018. As a response to the acquisition, the rest of the industry does more research effort, possibly to catch-up and to compensate for the disappearance of the independent startup.²⁴ Let us also note that these results contrast with Affeldt and Kesler 2021, who show that outsiders invest less in the product - in their context, an app - development after its acquisition by a tech giant.

²⁴For the non-merging party, we do not observe a difference between the citations to the big tech's core patents and the others.

8 Conclusion

With this paper, we aim to bring empirical evidence of the effect of big tech acquisitions on acquired innovative technologies. Information provided by the patent system allows us to track technologies before and after they are bought by these dominant firms. 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 are 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, we find no slowing down of the development of these Big-Tech acquired technologies by the rest of the industry, which means that their improvement potential has not been exhausted after acquisition. On this basis, we conclude that technology maturity cannot explain the slowing down in the development of Big Tech-acquired technologies. Instead, our analysis at the technology level indicates that a competitive motive could be driving this result; we find that the slowing down of the improvements made by Big Tech to acquired technologies is mainly observed for technologies in which they have focused their own innovative efforts. This last result could be driven by acquisitions strategies aiming to protect from the competitive threat that the target technology represents to its acquirer.

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 industries.

References

- 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.
- 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.
- 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.
- Allison, John R and Mark A Lemley (1998). “Empirical evidence on the validity of litigated patents”. AIPLA QJ 26, p. 185.
- Argentesi, Elena et al. (2021). “Merger policy in digital markets: an ex post assessment”. Journal of Competition Law & Economics 17.1, pp. 95–140.
- 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.
- Belleflamme, Paul and Martin Peitz (2015).
Industrial organization: markets and strategies. Cambridge University Press.
- Bourreau, Marc, Bruno Jullien, and Yassine Lefouili (2021). “Mergers and demand-enhancing innovation”. CEPR Discussion Paper No. DP16031.
- 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.

- Cabral, Luis (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.
- Caffarra, Cristina, Gregory Crawford, and Tommaso Valletti (2020). “How tech rolls’: Potential competition and ‘reverse’killer acquisitions”. Antitrust Chronicle 2.2, pp. 1–9.
- CMA (2021). “CMA directs Facebook to sell Giphy.” URL: <https://www.gov.uk/government/news/cma-directs-facebook-to-sell-giphy>.
- Cook, Tim (2019). “Berkshire Hathaway’s annual shareholder meeting.” Interview by Becky Quick on CNBC.
- Crémer, Jacques, Yves-Alexandre de Montjoye, and Heike Schweitzer (2019). “Competition policy for the digital era. Final report”. European Commission, DG Comp: Brussels.
- 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.
- 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”.
- 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.
- 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.
- 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.
- Gautier, Axel and Joe Lamesch (2021). “Mergers in the digital economy”.
Information Economics and Policy 54, p. 100890.
- 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.
- Haucap, Justus, Alexander Rasch, and Joel Stiebale (2019). “How mergers affect innovation: Theory and evidence”.
International Journal of Industrial Organization 63, pp. 283–325.
- Huntington-Klein, Nick (2021).
The effect: An introduction to research design and causality. Chapman and Hall/CRC.
- 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.

- 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.
- 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 Søgaaard (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.
- Lampe, Ryan (2012). “Strategic citation”. Review of Economics and Statistics 94.1, pp. 320-333.
- 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.
- 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/>.
- Moraga-González, José Luis, Evgenia Motchenkova, and Saish Nevrekar (2022). “Mergers and innovation portfolios”. The RAND Journal of Economics 53.4, pp. 641-677.

- 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.
- Parker, Geoffrey, Georgios Petropoulos, and Marshall Van Alstyne (2021). “Platform mergers and antitrust”. Industrial and Corporate Change 30.5, pp. 1307–1336.
- 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.
- 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.
- 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”.
- Simon, Phil and M Joel (2011). “The age of the platform: How Amazon, Apple, Facebook, and Google Have Redefined Business”. Las Vegas: Motion Publishing.
- “Stigler Report” (2019). Stigler Committee on Digital Platforms, Final Report.
- 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.
- “TomTom/Tele Atlas” (2008). Case M.4854, Commission decision of 14 May 2008, paragraphs 244–250.

Appendix I

Use of 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 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 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; 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.

Appendix II

Figure 8: Distribution of filing dates for all citing patents from the Patent Views database (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.

Appendix III

Inverse probability weighting

In order to make acquired and non-acquired patents comparable in all aspects 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-treatment 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 ($CitGR_{p,Pre}$):

$$P(A_p = 1 | CitGR_{p,Pre}, Cit_{p,Pre}) = \Phi(\alpha + \beta CitGR_{p,Pre}, Cit_{p,Pre}) \quad (8)$$

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.

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 5 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 3).

Table 5: Balance tables

Raw sample (before trimming and weighting)			
Variable	(1) Not acquired	(2) Acquired	(3) Acquired vs Not
$\overline{Cit_{pre}}$	0.232 (0.869)	0.294 (0.806)	0.062 (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 (0.758)	0.276 (0.735)	0.059 (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.

Appendix IV

Sharp event study - Identification

In a paper studying the impact of having children on the gender wage gap, Kleven, Landais, and Sogaard (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.):

$$Cit_{p,t} = f(J' \tau_j, x(J', z_{p,t}) \tau_x, z_{p,t} \tau_z) \quad (9)$$

where $J' = \sum_{j \neq 0} I_j = t$ 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 be 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_+):

$$E [Cit_{p,t_+} - Cit_{p,t_0}] = E[f(1, x(1, z_{p,t_+}), z_{p,t_+}) - E[f(0, x(0, z_{p,t_0}), z_{p,t_0})] \quad (10)$$

where the smoothness of the average citations path absent acquisition would imply that $E[F(0, x(0, z_{p,t_+}), z_{p,t_+})] \approx E[F(0, x(0, z_{p,t_0}), z_{p,t_0})]$. 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 ?? 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:

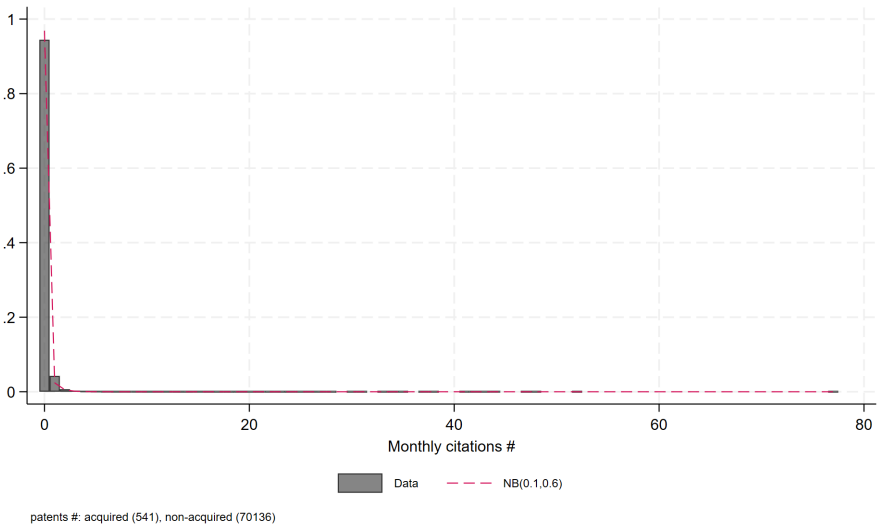
$$E [Cit_{p,t_{++}} - Cit_{p,t_0}] = E[f(T, x(T, z_{p,t_{++}}), z_{p,t_{++}}) - E[f(0, x(0, z_{p,t_0}), z_{p,t_0})] \quad (11)$$

The differences between this impact measure and equation 10 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, 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) might not properly capture the long-run impact of acquisition. We therefore propose with model 2 a second solution to capture long-term effects of acquisition, by using a control group to account for the citations trend absent acquisition.

Appendix V

Negative Binomial distribution of the citations count

Figure 9: 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.

Appendix VI

Robustness checks

Additional regressors

We rewrite model 1 to control for additional determinants of the number of citations received by a patent:

$$Cit_{p,j,t,d} = f(J'\theta^6, age_{p,d}\beta^6, M'\gamma^6, firm_j\xi^4, Z_p v^1) \quad (12)$$

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 GAFAM 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 and it acquires more core technologies than the others (see Table 2.1). 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 10 (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 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 10 (b).

The inclusion of these additional regressors seem to have little impact on our results, as the estimates presented on Figure 10 appear to be very similar to our baseline results presented on Figure 2.

Figure 10: Big Tech citations to acquired patents relative to acquisition

(a) Controlling for Microsoft FE (b) Controlling for technology fields and publisher's origin




Notes: The graph shows the incident rate ratios for acquired patents: $e^{\hat{\theta}_t^{\phi}}$ from model 12. These coefficients are estimated on a balanced sample of patents in a 4 year-window around acquisition.

Alternative study periods

Extending the pre-treatment period On Figures 11 and 12, we replicate the results from Models 1 and 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 extrapolation of the counterfactual can thus be based on the last portion of the pre-intervention period ($t = -3$ to $t = -1$).

Figure II: 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. These coefficients are estimated on a balanced sample of patents in a 4.5 year-window around acquisition.


Figure 12: 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: $e^{(\hat{\theta}_t^2 + \hat{\alpha}_t^1)} - e^{(\hat{\theta}_t^2)}$ from model 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 2016.²⁵ On the below figure, we observe that the evolution of citations just after acquisition follows a very similar trend to the baseline sample: citations increase significantly after acquisition. But, as the observation period is reduced, we do not capture the slowing down in citations, i.e. during 1.5 years, the effect of acquisition is positive.

Figure 13: 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. These coefficients are estimated on a balanced sample of patents in a 2 year-window around acquisition.

Including Motorola As an alternative check, we include *Motorola Mobility*, acquired by Google in August 2011 and later (January 2014) sold to Lenovo. Motorola was not

²⁵Because 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 859 Big Tech acquisitions from Table 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 859 Big Tech acquisitions).

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 1080 patents at acquisition among which 125 are cited by Google by June 2017, we propose an alternative study period that allows to include it.

Figure 14: Big Tech citations to acquired patents relative to acquisition



Notes: The graph shows the incident rate ratios for acquired patents: $e^{\theta_1^i}$ from model 1. These coefficients are estimated on a balanced sample of patents in a 3 year-window around acquisition.