Big Tech Acquisitions and Innovation: An Empirical Assessment

Laureen de Barsy[1*] and Axel Gautier[2*]

[*]HEC Liège, LCII, Université de Liège

March 14, 2023

PRELIMINARY VERSION

Abstract

Among the many challenges that poses the digital economy in the protection of competition, practitioners and academics have become increasingly concerned about the role of acquisitions by the main players of the sector: Alphabet, Apple, Meta, Amazon and Microsoft (often grouped under the label Big Tech). In particular, higher levels of concentration on this market can affect innovation. While the theoretical mechanisms explaining the effects of acquisitions on innovation in the digital sector are now relatively well understood, empirical evidence remains scarce.

This paper contributes to filling in the gap by looking into the fate that awaits patent-protected technologies acquired by Big Tech. We create a dataset of all public big tech acquisitions since 1996. For each of these acquisitions, we collect information on the acquired intellectual property. Our focus on patent data allows us to track the development steps of a technology as it moves across firms. Because 'prior art' is included in a patent by citations to previous patents, the evolution of the number of forward citations to acquired technologies around the time of acquisition can be used to study the impact of acquisition on the development of these

[1]ldebarsy@uliege.be
[2]agautier@uliege.be

technologies. We find that improvements made by Big Tech to acquired technologies slow down after acquisition, and that this cannot be associated with a 'normal' trend in the development of an acquired technology.

1 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 label 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 now represent the most valuable brands 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 poses the digital economy in that regard (e.g. strong network effects, multi-sidedness, data-driven economies of scope, etc.), the role of mergers and acquisitions (M&A) by Big Tech is increasingly considered.[1] In an interview on CNBC, Tim Cook (2019), Apple's CEO, highlighted the very high rate at which the platform acquires start-ups: *"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 the very high rate at which Big Tech acquires smaller firms, very few of these acquisitions are reviewed by a competition authority[2] 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 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).

The first concern about digital M&As is that they can lead to a loss in consumer welfare due to enhanced market power. In a static environment, economic theory predicts that, by relaxing a competitive pressure, horizontal mergers necessarily lead to higher prices. But their effect on innovation (and thus on future prices and products quality) is also key and it can be used as an argument in the "balance of harm" approach of competition authorities. As such, three different dimensions of innovation can be

[1]See for instance Argentesi et al. 2021; Crémer, Montjoye, and Schweitzer 2019; "Stigler Report" 2019.
[2]More than 97% of acquisitions in the technology sector have reportedly never been subject to scrutiny by a competition authority (Kwoka and Valletti 2020).

considered: innovation by the acquired start-up, by the acquirer's competitors and by the acquirer itself.

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, Moraga-González, Motchenkova, and Dijk 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).

Digital M&A might also impact innovation by the acquirer's competitors, actual or potential. 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.

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). A merger increases appropriability since the pioneering firm now also benefits from these positive externalities reaching the other merging party (Federico, Langus, and Valletti 2018). 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). The start-up might also not have the resources to bring the project to the market, in which case acquisition can foster innovation if the platform has an incentive to develop it (Motta and Peitz 2021, Fumagalli, Motta, and Tarantino 2020). The main driver of the negative effects of M&A on innovation is their impact on the market structure. Innovation is a competitive tool through which a firm can steal business from its competitors. By merging, previously competing firms internalise these business stealing effects, which thus reduces their incentives to innovate (Federico, Langus, and Valletti 2018; Federico, Morton, and Shapiro 2020; Motta and Tarantino 2016). A second mechanism through which M&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). Finally, incumbents might use acquisitions as a way to get rid of start-ups that represent potential competition because they are developing substitute products to their own (Cunningham, Ederer, and Ma 2021, Fumagalli, Motta, and Tarantino 2020). Cunningham, Ederer, and Ma 2021 document 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. the acquisition kills the innovation.

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 recognised 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. In the digital context, the literature has so far adopted two main approaches to track development activity after a big tech acquisition: i. whether the company website reveals that the acquired product is still being supplied, maintained or upgraded (Gautier and Lamesch 2021) and, ii. in the specific case where the acquired product is a mobile application, whether it is still undergoing functional updates (Affeldt and Kesler 2021). According to these studies, 50% to 60% of products are discontinued after a big tech acquisition. 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 or be integrated in a new product. Little is known about the development of technologies after acquisition and this paper intends to fill in this gap.

To assess the impact of big tech acquisitions on innovation, and instead of tracking project-level development activities, this paper focuses on the projects' innovative part, which is protected by patents. By tracking patents as they move across firms, we are able to identify whether the technology is still being developed after acquisition. More specifically, the patent system is such that, when some inventor builds 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. Using the time series nature of our data, we develop a methodology to isolate the effect of acquisition on Big Tech citations to acquired patents from the effect of (observed or unobserved) time trends. First, 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. By means of a propensity score weighting design, we then compare the remaining time trends in Big Tech citations to acquired patents with respect to comparable non-acquired patents. Our empirical analysis shows that, after acquisition, the developments made by Big Tech to acquired technologies slow down. However, this trend in the acquired technology development is not observed among other firms than their acquirer; these firms cite Big Tech-acquired patents steadily around 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.

In Section 2, we describe the main features of Big Tech acquired technologies (2.1) and the construction of our working datasets (2.2). Sections 3 discusses our empirical strategy to take out the effect of endogenous factors from the evolution of Big Tech's citations to acquired patents around the time of acquisition. Our baseline results are presented in Section 4.1, while Section 4.2 proposes some additional analyses to try and interpret these results. Section 5 concludes.

2 Empirical Methodology

For our analysis, we construct 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

We first create a dataset of firm's acquisitions by Alphabet, Amazon, Apple, Meta and Microsoft based on the Standard & Poor's CapIQ database.[3] We retrieve information on the identities of the acquired firms and on the dates at which their acquisitions were announced and closed. Then, we use the OECD patent statistics, built on the PATSTAT database, to match acquired firms with intellectual property; we obtain a sample of all published patents filed by a Big Tech-acquired firm to the European Patent Office (EPO), the US Patent and Trademarks Office (USPTO) or through the Patent Cooperation Treaty (PCT).[4]

Among all 707 public big tech acquisitions closed between January 1996 and January 2021, we identify 275 firms that own at least one patent (or that have filed at least one patent application) before being acquired. Table 1 presents the steps in the construction of the working dataset. Since we identify technology developments by tracking patents as they move across firms, we will restrict our analysis to those 275 acquisitions associated with patent-protected technologies.

[3] Let us note that acquisitions of very small firms might not be covered by the CapIQ database.
[4] Patents published under the target's name after acquisition are considered as published by the acquirer.

Table 1: Number of Big Tech acquired firms

	Firms acquired by Big Tech btw. 1996 and 2021 (CapIQ data)	Acquired firms with at least one published patent (OECD data)	Acquired firms with at least one patent filed pre-acquisition
Amazon	92	40	29
Apple	93	51	49
Facebook	86	26	23
Google	229	89	78
Microsoft	207	103	96
TOTAL	707	309	275

Note: This table illustrates the steps that are taken to select, among all Big Tech-acquired firms, those that have patented a technology.

2.1.2 Patent data

We collect information on the patents acquired by Big Tech through the acquisition of the company that filed these patents. Patents are included in our database irrespective of whether they are further granted by the patent office to which they were filed. They are identified based on their application number.

To control for potential trends in the technology development over a patent's life, we retrieve information on the patent age (based on its publication date) for all published patents filed at the EPO, PCT and USPTO.[5] For granted USPTO patents, we also observe the following information on the patent's technological and economic value: patent scope, family size, grant lag, number of claims, generality, originality, and radicalness. Based on these 7 indicators, we construct a patent quality index. The definitions of these variables and the reasoning behind their inclusion in the quality index are presented in Appendix I. All these variables are normalized such as to be centered in 0, with higher values associated with higher quality levels, and their average defines the composite quality index. This means that a patent associated with a positive quality index can be considered as more valuable than average, and a patent associated with a negative quality index can be considered as less valuable than average. As could be expected, we observe that Big Tech-acquired patents are more valuable than the average patent (see Appendix II). As an alternative to this quality index, we will use one vector that captures as much as possible of the variation in the data along the seven indicators of a patent's value. More specifically, we intend to approximate our 7-D quality space by a linear combination of all 7 (normalized) indicators along which the spread of patents is maximised. This vector is computed using a principal components analysis,

[5] Information is available for patents published from 1978 onwards.

as described in Appendix III .

2.1.3 Citations data

The use and the further development of a patented technology can be proxied by the 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.

To retrieve that information from the OECD database,[6] we use the application identifiers of all patents containing a citation to a patent filed by a Big Tech-acquired firm. Patents cited by their acquirer can then be identified by matching these application identifiers to the filing firms. In addition, we can estimate the date at which each citing patent was filed; 18 months before publication for EPO and PCT patents, and 9 months before publication for USPTO patents. [7] On this basis, we can derive the number of citations received by a given patent at a monthly level. Let us note that, while *citing* patents information is available for EPO and PCT patents irrespective of whether they are further granted, it is only available for USPTO granted patents. Because the citations data is available until July 2021, and to avoid biases due to some citing US patents not yet being granted by that time (and hence not observed), we end our study period 55 months before the data collection, in December 2016. This buffer period is defined based on the observation that 90% of US patents are granted within 55 months of their application.[8]

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 acquisi-

[6]The OECD Citations Database covers all citations present in EPO and PCT patent documents published from 1978 onwards and all citations made in USPTO patent grants from 1976.

[7]We do not directly observe a patent filing date. 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, respectively, 100% and 95% of all EU and US patent applications (Tegernsee 2012), and 97% of all PCT patents are published within two weeks of this limit (WIPO 2022). While this is a minimum delay for EPO and PCT patents, earlier publication is often observed at the USPTO: half of US patent applications are published within 9 months after they were filed (Martin 2015).

[8]Author's calculations based on the 'grant lag' from the OECD Patent Quality Indicators database, July 2021.

tion 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, for instance because they are simply not patentable (Belleflamme and Peitz 2015), or due to a low probability of imitation and/or high costs of patenting (e.g. hiring patent specialists to prepare the application, paying the filing administrative costs and the renewal fees). However, patent data allows us to capture most significant Big Tech acquisitions, as 70 out of the 100 biggest acquired firms (i.e. with a total funding above $2.5 million)[9] had patented some invention and hence are included in the patent database.

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.

[9] Funding data is retrieved from Crunchbase.

2.1.4 Summary statistics

In the end, for each patent in our database, we can identify the acquirer, the timing of acquisition (announced date and closed date), the patent's age, the patent's quality, and the number of forward citations made every month to this patent. We construct a dataset containing all patents belonging to one of the 275 Big Tech-acquired firms, and we select those firms that have published, pre-acquisition, at least one patent further cited by their acquirer; approximately 900 acquired patents are never cited by the acquirer. We end up with a working sample of 187 firms, i.e. 187 patent portfolios. Table 2 presents summary statistics of these data samples at the patent portfolio level. On average, portfolios that are cited by their acquirer are more likely to have been published at the USPTO[10] than portfolios that contain no patents cited by their acquirer.

Table 2: Big Tech acquired patents portfolios

		Firm size ($mm)		Portfolio size (patents #)		Patent age at acquisition		High quality patents	US-published patents
	Count	Mean	SD	Mean	SD	Mean	SD	%	%
Big Tech acquired portfolios									
AMZN	29	399	448	26	78	4,73	5,40	74%	64%
APPL	49	187	243	71	221	4,00	2,39	67%	66%
FCBK	23	7.476	10.994	14	38	3,73	4,13	43%	52%
GOOG	78	1.282	2.544	48	252	4,29	3,42	80%	61%
MSFT	96	1.972	5.585	27	83	3,61	2,80	62%	59%
TOTAL	275	1.384	4.132	40	172	4,01	3,41	70%	62%
Big Tech acquired portfolios cited by their acquirer									
AMZN	15	326	295	18	27	3,31	2,25	74%	70%
APPL	36	215	256	95	254	4,55	2,70	79%	88%
FCBK	10	11.139	12.697	27	56	3,96	3,64	45%	81%
GOOG	52	1.412	2.838	70	307	3,94	2,24	84%	84%
MSFT	74	2.336	6.071	34	93	4,01	2,84	65%	79%
TOTAL	187	1.664	4.639	54	206	4,03	2,64	75%	82%

Notes: This table provides summary statistics on Big Tech-acquired patents portfolios. High quality patents are defined as those associated with a quality index superior to the median quality index.

2.2 Working sample

[10] Based on a probit regression at the 1% level.

2.2.1 Acquired and non-acquired patents

We consider patents that are acquired by a tech giant and that receive at least one forward citation (further simply referred to as 'acquired patents'). If citations are made by the same citing patent to cited patents belonging to a same patent family, we keep the cited patent with the earliest publication date. Next, we select a balanced panel of Big Tech-acquired patents observed every month between 1 year before their acquisition (closing) date and 3.5 years after,[11] such that our time series cover a period of 4.5 years. 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: 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.[12] 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 before simulated acquisition and 3.5 years after. On this basis, we obtain two groups: i. a balanced panel of patents acquired between 1996 and 2021 and observed in a 4.5 year-window around acquisition, ii. a balanced panel of patents that were never acquired, but that have been assigned a placebo acquisition date between 1996 and 2021 and are observed in a 4.5 year-window around this placebo. Table 3 presents the number of patents in these two groups: 844 patents undergo an acquisition event, and 103,198 are assigned a placebo acquisition date.

[11] We assume citations are observed from 6 months before the patent publication, because 99% citations in our sample are made less than 6 months before the patent publication.

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

Table 3: Patents observed over the whole study period

	Big Tech acquired	Big Tech non-acquired
Amazon	79	8,321
Apple	122	28,169
Facebook	1	4,151
Google	338	23,764
Microsoft	304	38,793
TOTAL	844	103,198

Note: This table presents the number of observations contained in the balanced sample of patents observed in a 4.5 year-window around (simulated) acquisition. There are two reasons why Facebook is underepresented. First, the company is not very active from a patenting point of view. Second, Facebook started acquiring smaller firms later than the other tech giants, so most of its patented acquired technologies are not observed 3.5 years after acquisition.

2.2.2 Inverse probability weighting

Before comparing the development of Big Tech-acquired and non-acquired patents, we must ensure that the two groups are comparable in all aspects except for their acquisition status, as if acquisition had been fully randomized. In particular, we must deal with a selection bias.

Acquired patents are acquired for a reason; the acquirer must see some potential in them. In order to obtain two comparable groups, 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). A patent's quality is an obvious determinant of both its acquisition status and the citations it receives. Some other determinants are unobserved. They will be captured by including the number of citations received by a given patent before acquisition. However, the regressors included in the propensity scores must be exogenous to the acquisition event. Because most acquisitions are announced before they are actually closed, this announcement might already impact the citations to the to-be-acquired patents. This means that the event might have some anticipation effects. Therefore, we aim to use a period of time during which the announcement has not taken place yet. From our working sample, we observe that only 2% Big Tech acquisitions are announced more than 8 months before their closing date. On this basis, we choose to use the number of citations up to 9 months before acquisition (i.e. during the first three months of our time series).

To obtain the propensity scores, 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, its quality ($Qual_p$) and the number of citations it receives during the first three months of our time series ($Cit_p(0)$):

$$P\left(A_p = 1|Qual_p, Cit_p(0)\right) = \Phi(\alpha + \beta Qual_p + \gamma_1 Cit_{p,-11} + \gamma_2 Cit_{p,-10} + \gamma_3 Cit_{p,-9}) \qquad (1)$$

where 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 covariate values have a positive probability of being both acquired and non-acquired. In other words, conditional on these scores, the distribution of the endogenous covariates should be similar between acquired and non-acquired patents (Austin 2011). For instance, because we observe that patents with a higher quality are more likely to have been acquired,[13] the propensity score associated with a high quality patent is higher.

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. On this basis, we end up with a sample of 663 acquired patents (all acquired patents for which we have information on quality) and 69,396 non-acquired patents. This represents, respectively, 79% and 67% of the original balanced sample presented in Table 3. Because patent quality is not observed for EP- and WO-published patents and pending US-published patents, only US-granted patents are included in this trimmed sample. 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 in-

[13] Based on a t-test at the 1% level on the balanced sample, Big Tech-acquired patents are found to be statistically higher quality (mean index = 0.05) than non-acquired patents (mean index = 0.02).

verse 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". Balance is the assumption that, after weighting, there are no more meaningful differences between acquired and non-acquired patents in variables included as regressors to compute the propensity scores. This ensures that the inverse probability weighting is appropriate to close the propensity scores channel through which these regressors 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. The balance tables in Appendix IV present the results of this test before and after applying the inverse probability weighting. We observe that, compared to the raw sample, the differences in means between acquired and non-acquired patents in the new trimmed and weighted sample are reduced. However, in the end, individual variables do not matter; the only one variable used to weight observations is the propensity score itself. So, for our weighted sample to be properly balanced, the propensity score distributions for the two patent groups should be nearly identical. We present in Appendix V the overlaid density plots of the obtained propensity scores. Because the propensity score distributions for the treated group of acquired patents and the control group of non-acquired patents appear to be quite similar, we conclude that, after dropping observations outside the range of common support and weighing observations based on their inverse probabilities, these two patent groups will become comparable.[14]

[14]We also see from Appendix V that the propensity scores take very small values. This is due to the large difference between the sample sizes of acquired and non-acquired patents; since there are much fewer acquired patents, each of them is given a relatively higher weight. In the next section, we will test whether our results are robust to selecting two patents groups that are similar in size.

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. Let us highlight that our analysis does not aim to compare pre- and post-acquisition levels in citations, as pre-acquisition citations levels might not be exogenous to the acquisition event. Instead, we will try and identify the changes in citations that can be attributed to acquisition over event time.

3.1 Technology developments around the time of acquisition

Our model is defined in two stages. First, we formalize the relationship between the acquisition event time and the development of the acquired technology by the acquiring firm as measured by citations to the associated patents. Second, we design a method to take out the effect of factors endogenous to the acquisition event that could explain the trend in citations.

3.1.1 Observed time trends

We study the evolution of the number of forward citations by the acquirer as a function of event time dummies, which represent the months in which citing patents are filed with respect to the time of acquisition $t = 0$. 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. Two main determinants should be accounted for: life-cycle trends (i.e. the number of forward citations might depend on the stage of a patent's life) and business-cycle trends (i.e. the industry's R&D might be more or less dynamic in given years).

We denote by $Cit_{p,t,d}$ the number of forward citations by the acquiring firm to patent p at event time t and date d. 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 date m dummies. 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:

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

where J' is a vector containing the time dummies ($t = -11, -10, ..., -1, 0, 1, ..., 41, 42$) excluding the base category $t = 0$, and M' is a vector containing the date dummies ($d = 1996m1, 1996m2, ..., 2016m12$). To define the function $f(.)$, we must account for the nature and distribution of the response variable: the citations count. The most widely used model for a count regression is the Poisson distribution. However, the Poisson model assumes that the mean and variance of the errors are equal. In our case, the variance of the citations count is much larger than its mean: a majority of patents in the data set are only cited once, but a few patents are cited many times (see Appendix VI). 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). The negative binomial regression model can be written as:

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

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

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. This is estimated by the event time impact on the number of citations: $\hat{\theta}_t^1$.

3.1.2 Unobserved time trends

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). We rewrite model 2 as follows:

$$Cit_{p,t,d} = f(J'\theta^2, A_p\iota^1, J'A_p\alpha^1, age_{p,d}\beta^2, M'\gamma^2) \tag{3}$$

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

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: $\widehat{\alpha}_t^1$.

4 Results

We present below estimates of the impact of a big tech acquisition on the development of the acquired technology by the acquiring firm as measured by citations to the associated patents. Model 2 is estimated on the balanced panel of Big Tech-acquired patents, and model 3 is estimated on the balanced panel of trimmed Big Tech-acquired and non-acquired patents where the contribution of each observation has been multiplied by its weight (i.e. its inverse probability as described in Section 2.2).

4.1 Baseline sample

4.1.1 Impact of Big Tech acquisitions on the acquired technology development by the acquirer

Figure 1 plots the estimated event coefficients from models 2 and 3 across event time. As defined above, these are 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 ($\hat{\theta}_t^1$) for acquired patents with respect to non-acquired patents ($\widehat{\alpha}_t^1$). The figure includes 95% confidence bands around the event coefficients.

We see that, before acquisition, Big Tech citations to soon-acquired patents grow faster than for non-acquired patents. Since patents citations are used as a measure of technology improvements, this pre-acquisition observed positive trend in citations shows an increasing interest by the tech giant for the acquired technology. This interest seems to fade away around 6 months after acquisition as improvements made to the technology by its acquirer start slowing down. From around 2.5 years after acquisition, we even observe a drop in the number of citations to acquired patents with respect to non-acquired patents ($\hat{\alpha}^1_{t>30}$ take negative values).

Two main interpretations can be put forward to explain these results. First, the observed citations pattern might just illustrate a normal trend in an acquired technology development; the technology might be improved until it reaches maturity, and this might coincide with the acquisition event (we will further refer to this interpretation as 'technology maturity'). Second, Big Tech might be shifting its innovative efforts away from acquired technologies. In Section 4.2, we propose to further investigate the credibility of these two potential interpretations.

Figure 1: Big Tech citations to acquired patents relative to the time of acquisition

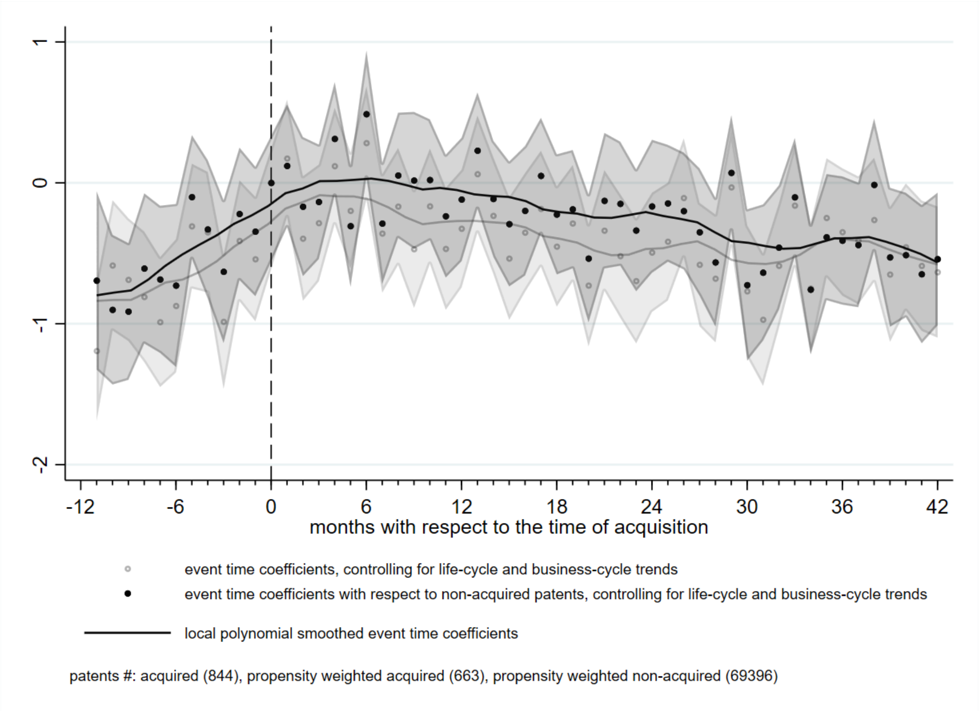

Notes: The graphs show the event time coefficients: θ^1 from model 2 (grey dots), and α^1 from model 3 (black dots). These coefficients are estimated on a balanced sample of patents in a 5 year-window around (simulated) acquisition. The shaded bands represent 95% confidence intervals.

4.1.2 Robustness checks

To test for the robustness of our results, we propose in Appendix VII to replicate our analysis based on four alternative samples: i. one from which outlier observations are excluded, ii. one that reduces the study period from 4.5 years to 2 years around acquisition such as to capture more Big Tech acquisitions, iii. one with more equally sized groups of acquired and non-acquired patents, and iv. one with probability weights computed using the principal component as an alternative to our quality index (as described in Appendix III). By estimating models 2 and 3 based on these four alternative samples, we find similar event time coefficients to those obtained using the baseline sample.

Next, because we observe that most Big Tech acquisitions are associated with a time lag between their announcement and closing date, we replicate our analysis by using the announcement date (instead of the closing date) as the event time. The event time estimates obtained from this robustness check are presented in Appendix VIII (a). Because most acquisition announcements take place before the closing dates, part of the pre-closing positive trend in the acquired technology development takes place after the announcement event ($\hat{\theta}^1_{3<t<10}$ take positive values). But our main result holds; after acquisition, acquired technologies appear underdeveloped compared to non-acquired technologies ($\hat{\alpha}^1_{t>0}$ take zero or negative values).

Finally, 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 2 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. If we were to ignore this over-dispersion of our data and constrained δ to zero, we would obtain the event time estimates from a Poisson model, as presented in Appendix VIII (b). Based on these (biased) estimates, we find a slight positive evolution in citations until 6 months after acquisition ($\hat{\alpha}^1_{t=4}$ is positive), which then becomes negative just like in the baseline model ($\hat{\alpha}^1_{t>6}$ take zero or negative values).

4.2 Exploring the 'technology maturity' interpretation

4.2.1 Technology developments by other firms than the acquirer

We discussed above that the slowing down of the acquired technology development might be explained by 'technology maturity'. This would for instance be the case if, before acquisition, the incumbent platform and the start-up would be working in parallel on similar R&D projects. In order to strengthen its own innovation effort, the platform might decide to buy the start-up. Thanks to the pooling of skills and assets from both companies, the technology would then reach its maturity. After this point, each marginal increase in the innovative effort would result in smaller improvements of the technology, which would explain that it is less developed after acquisition. To further explore this potential interpretation of our results, we will look at the acquired technology developments by other firms than the acquirer. If the technology is indeed reaching maturity short after acquisition, we should observe a slowing down of its development not only by the acquirer, but by the industry as a whole.

We refine model 2 by including, as additional regressors, fixed effects that are constant for all cited patents published by a same acquired firm ($firm_j$):

$$Cit_{j,p,t,d} = f(J'\theta^3, age_{p,d}\beta^3, M'\gamma^3, firm_j\xi^1) \tag{4}$$

Models 2 and 4 are estimated on two separate samples: balanced panels of Big Tech-acquired patents cited at least once over our study period by 1) the acquirer and 2) other firms than the acquirer.[15] The estimated event time coefficients ($\hat{\theta}_t^1$ and $\hat{\theta}_t^3$) are presented on Figure 2, separately for these two citing groups. As can be observed, and contrary to citations by the acquirer (a), citations by other firms than the acquirer (b) are relatively constant around the time of acquisition and even slightly grow from around 2 years after acquisition.[16]

These findings are not in line with the 'technology maturity' interpretation, according to which the slowing down of an acquired technology development after acquisition would be explained by diminishing returns to the innovative effort. Instead, we find no negative effect of acquisition on citations by the rest of the industry, which suggests that the improvement potential of the technology has not been exhausted after acquisition. Therefore, we argue that the observed acquirer's citations pattern illustrates a tendency on the part of Big Tech to under-develop acquired technologies.

[15] The data set on citations by other firms than the acquirer is constructed in the exact same way as for citations by the acquirer.

[16] A similar pattern is observe for the citations by the other Big Tech than the acquirer.

Figure 2: Citations to Big Tech-acquired patents relative to the time of acquisition

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

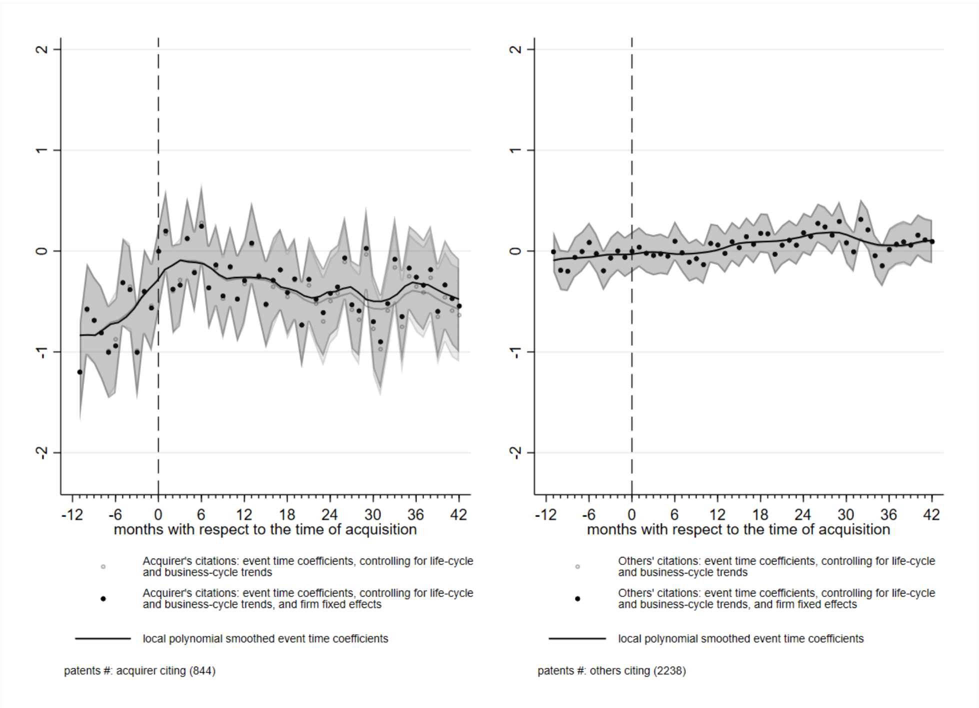

Notes: The graphs show the event time coefficients: θ^1 from model 2 (grey dots), and θ^3 from model 4 (black dots). These coefficients are estimated on a balanced sample of patents cited by their acquirer (on the left) and by other firms (on the right) in a 5 year-window around acquisition. The shaded bands represent 95% confidence intervals.

5 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. After taking out the effects of life-cycle and business-cycle trends, we find that Big Tech's citations to acquired patents with respect to non-acquired patents slow down after acquisition. 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 that citations to Big-Tech acquired patents by other firms than the acquirer are relatively constant around the time of acquisition, or even evolve rather positively after acquisition, which means that the improvement potential of the technology has not been exhausted after acquisition. On this basis, we conclude that the slowing down in the development of Big Tech-acquired technologies cannot be interpreted as a normal trend in a technology development and that, instead, Big Tech appears to be shifting its innovative efforts away from acquired technologies.

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.

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.

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

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.

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

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

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

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

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

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

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

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

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

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.

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

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

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

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

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

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

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

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.

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.

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

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

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

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 Søgaard (2019). "Children and gender inequality: Evidence from Denmark". American Economic Journal: Applied Economics 11.4, pp. 181–209.

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.

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

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

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

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

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

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

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

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

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.

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

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.

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

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.

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

"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.

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

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

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

Appendix I

Components of the patent quality index

As discussed in Section 2.1.2, we construct a patent quality index using indicators measuring patents technological and economic value: patent scope, family size, grant lag, number of claims, generality, originality, and radicalness. For each of these variables, we explain here, based on the OECD note accompanying the related data (Squicciarini, Dernis, and Criscuolo 2013), why they are good measures of a patent's quality.

Patent scope

The patent scope variable is constructed based on the number of distinct 4-digit IPC technology subclasses listed in the patent document. It has been shown that technologically broader patents are associated with more valuable firms (Lerner 1994).

Family size

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

Grant lag

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

Adjusted number of claims

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

Generality

The patent generality variable is constructed based on the number of distinct technology fields to which citing patents belong (adjusting for the total number of citing patents). The wider the range of fields, the more relevant the cited patent has been for subsequent innovation.

Originality

The patent originality variable is constructed based on the number of distinct technology fields to which cited patents belong. The broader the technology fields on which a patent relies, the more original the resulting innovation is expected to be (Trajtenberg, Henderson, and Jaffe 1997).

Radicalness

The patent radicalness variable is constructed based on the number of IPC technology classes listed in the cited patents documents, but in which the patent itself is not classified. The more diversified the array of technologies on which the patent relies upon, the more the invention should be considered radical (Shane 2001).

Appendix II

Distribution of the patent quality index

Figure 3: Patent quality index distribution (Density)

(a) all patents (b) Big Tech-acquired patents

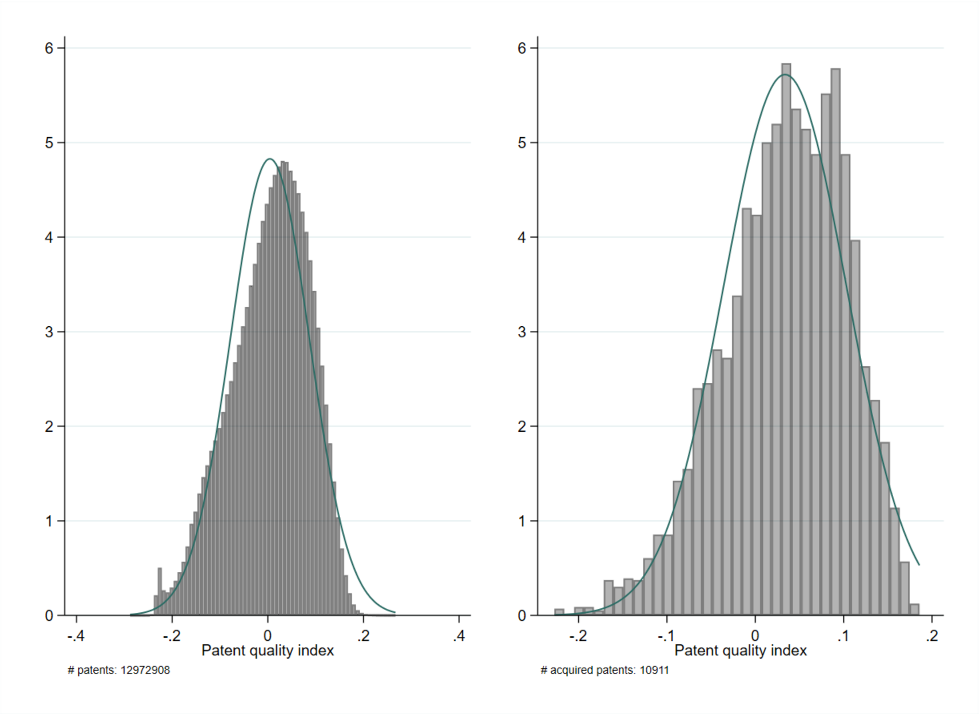

Notes: This graph displays the distribution of the patent quality index for all patents and Big Tech-acquired patents. The overlaying curve represents a normal density.

Appendix III

Principal Component Analysis

As an alternative to our patent quality index, we propose to define one vector that approximates the information contained in the seven indicators of a patent's value. To do so, we follow a principal component analysis. First, we center our observations around a point that has as coordinates the average values of all 7 (normalized) indicators variables. This point is called the center of the data. Second, we calculate the principal components. Principal component 1 ($pc1$) is calculated as the line that best fits the data while going through the center. Principal component 2 ($pc2$) is calculated as the next best fitting line that goes through the center and is perpendicular to $pc1$. The other components (pc_i, $i \in [3, 7]$) are computed in the same way, i.e. the best fitting lines that go through the center and are perpendicular to all $pc_{k<i}$. Third, we project our data onto these 7 axes to obtain the values of the principal components in our sample. The proportion of the variation in the data that each principal component accounts for can be computed as the distance between these data projections and the center of the data. This distance is called the eigenvalue of a component. To choose the number of components, we use the scree test developed by Cattell (1966), which assumes that, when the eigenvalues are plotted as a function of the number of components, the eigenvalues should decline gradually as more and more noise is modelled (Bro and Smilde 2014). In our data, the point of inflexion is located at the second component, so we retain $pc1$ and $pc2$. As we can see from Table 4, these two components accounts for 47% of the variation in the data. The importance of each of the 7 indicators of a patent's value for the obtained principal components is captured by the value of its eigenvector, presented in Table 5. The main determinants of $pc1$ appear to be the patent's generality and originality, i.e. the (normalized) number of distinct technology fields to which citing and cited patents belong, and $pc2$ is mainly determined by the patent's scope and family size, i.e. the (normalized) number of listed technologies and of countries where the invention is protected.

Table 4: Principal components

	Eigenvalue	Proportion (cumulative)
$pc1$	2.01039	0.2872
$pc2$	1.25212	0.4661
$pc3$	1.04398	0.6152
$pc4$	.935372	0.7488
$pc5$	.801907	0.8634
$pc6$	.586627	0.9472
$pc7$	.369598	1.0000
Patents #	5,453,200	

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

Table 5: Eigenvectors

	$pc1$	$pc2$	$pc3$	$pc4$	$pc5$	$pc6$	$pc7$
Patent scope	0.3870	0.5340	0.1108	0.0062	−0.4560	0.4003	0.4296
Family size	0.1659	0.5870	−0.1245	0.2564	0.7309	−0.1083	0.0260
Grant speed	−0.2148	−0.0588	0.4104	0.8719	−0.1303	0.0558	−0.0421
Adjusted claims	−0.1803	0.1340	0.7995	−0.3980	0.2515	0.2463	−0.1672
Generality	0.4969	0.0844	0.3430	−0.0269	−0.2160	−0.7407	−0.1798
Originality	0.5740	−0.1877	−0.0758	0.1131	0.0588	0.4611	−0.6331
Radicalness	0.4096	−0.5536	0.2029	0.0465	0.3570	0.0558	0.5933
Patents #	5,453,200						

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

Appendix IV

Balance tables

Tables 6 and 7 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 variables included as regressors to compute the propensity scores, for control observations (non-acquired patents) and treated observations (acquired patents) respectively. In the third column, we regress these variables on the observation's treatment value (acquired or not) to compute the differences of means and the associated standard errors.

Table 6: Raw sample (before trimming and weighting)

Variable	(1) Not acquired	(2) Acquired	(3) Acquired vs Not
Qual	0.017	0.052	0.035***
	(0.075)	(0.064)	(0.003)
Cit_{-11}	0.053	0.053	0.000
	(0.343)	(0.314)	(0.013)
Cit_{-10}	0.051	0.107	0.056***
	(0.316)	(0.800)	(0.013)
Cit_{-9}	0.053	0.163	0.110***
	(0.350)	(1.273)	(0.014)
Observations	69,495	663	70,158

Standard errors in parentheses.
Significance at the *** 1% level, ** 5% level, * 10% level.

Table 7: Working sample (after trimming and weighting)

Variable	(1) Not acquired	(2) Acquired	(3) Acquired vs Not
Qual	0.017	0.015	-0.002***
	(0.075)	(0.080)	(0.000)
Cit_{-11}	0.052	0.048	-0.004**
	(0.344)	(0.321)	(0.002)
Cit_{-10}	0.051	0.033	-0.018***
	(0.318)	(0.249)	(0.002)
Cit_{-9}	0.054	0.041	-0.012***
	(0.351)	(0.320)	(0.002)
Observations	69,396	663	70,059

Standard errors in parentheses.
Significance at the *** 1% level, ** 5% level, * 10% level.

Appendix V

Propensity score distributions

Figure 4: Propensity score distribution by acquisition status (Density)

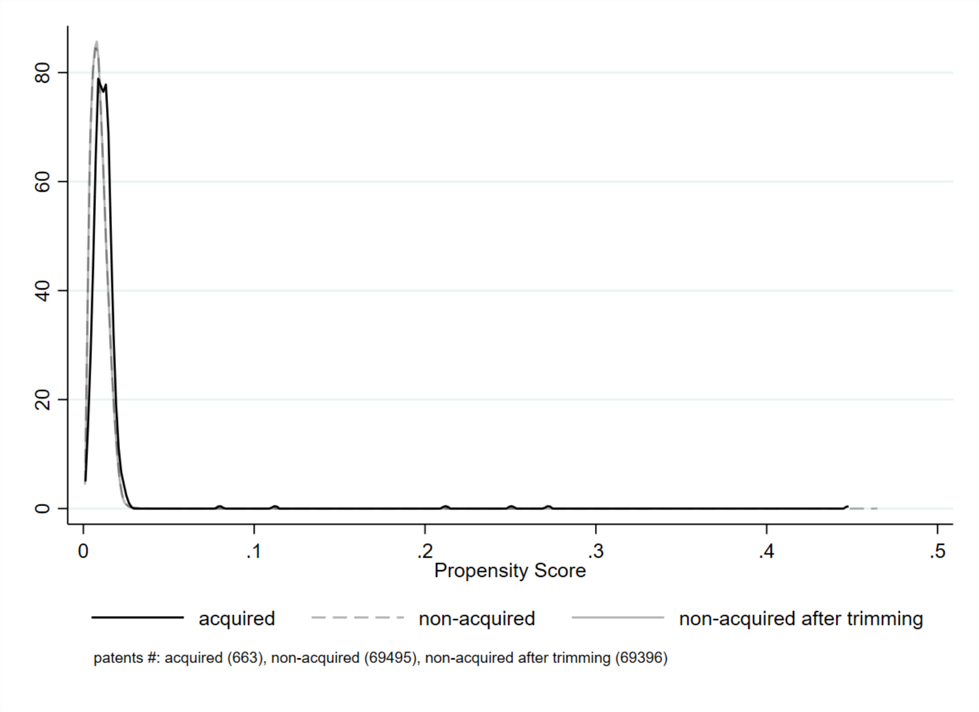

Notes: This figure shows the distribution of the propensity scores for acquired and non-acquired patents.

Appendix VI

Negative Binomial distribution of the citations count

Figure 5: Distribution of the count of Big Tech citations (Percent)

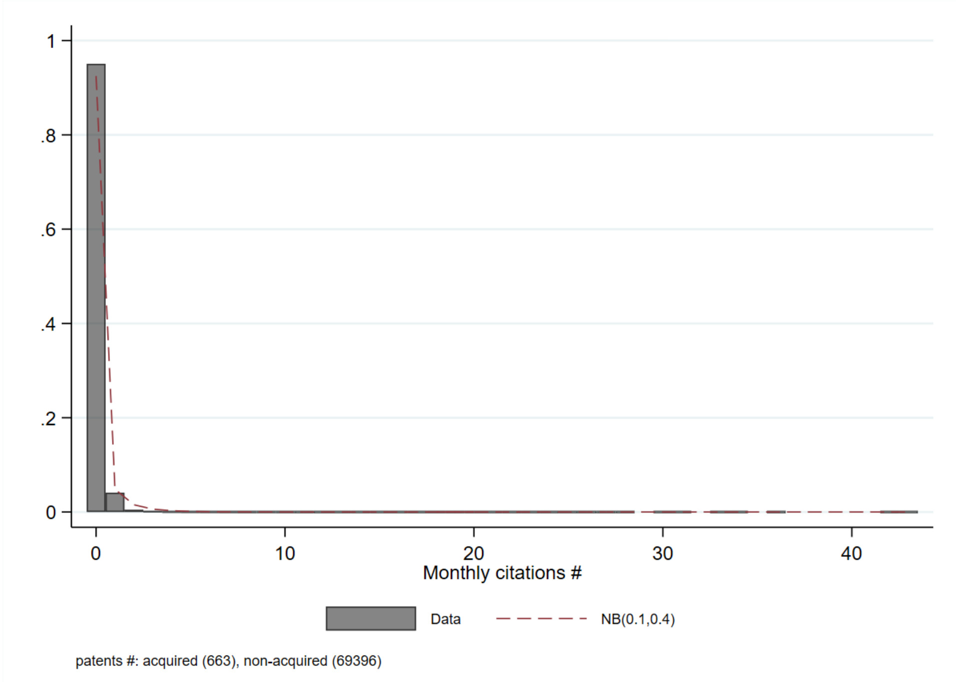

Notes: This figure shows an histogram of the number of citations received by a given patent in a given month, overlaid with a negative binomial density with the same parameters.

Appendix VII

Alternative data samples

As robustness checks, we propose to replicate our analysis based on four alternative samples. First, we exclude outliers by assuming citations follow a negative binomial distribution; patents associated with citations levels with an expected frequency of less than 0.5 are excluded. In practice, this implies that we only keep acquired patents that are cited less than 8 times in each time period. Estimating models 2 and 3 based on this alternative sample brings similar results to those obtained using the baseline sample: as can be observed from Figure 6 (a), acquirer's citations to acquired patents are slowing down after acquisition, and experience a drop in levels from around 2.5 years after acquisition. Second, we consider a shorter time series; from a period of 4.5 years around acquisition (1 year before and 3.5 years after), we reduce our study period to 2 years around acquisition (1 years before and 1 year after). This allows to capture more Big Tech acquisitions; because we end our study period in December 2016 to avoid biases in the citations count, restricting our sample to patents observed up to 3.5 years after acquisition meant that we could only use acquisitions undertaken until June 2013, which represent only half of all 707 Big Tech acquisitions from Table 1. By adding patents observed between 1 and 3.5 years after acquisition, we capture acquisitions up to December 2015, which represent 3/4 of Big Tech acquisitions. On Figure 6 (b), we can see that the drop in the number of citations is not yet significant by the first year after acquisition but we do observe, just like for the baseline sample, a positive trend before acquisition and a negative trend after acquisition.

A third working sample is defined in order to test whether the accuracy of our results is affected by the large difference between the sample sizes of acquired and non-acquired patents in the baseline sample (leading to low values of the propensity scores, see Appendix V). We intend to use all 663 acquired patents and 1,000 non-acquired patents. To do so, we first randomly draw (without replacement) 1,000 observations from the sample of non-acquired patents. We then compute the propensity scores based on the full acquired sample and on the 1,000 non-acquired random draws. As can be observed from Figure 7, the obtained propensity scores are larger than those from the baseline sample. After applying inverse probability weighting on the new samples, we estimate model 2 and we save the obtained coefficients estimates. This entire process is repeated 100 times. In the end, we obtain 100 coefficients estimates, which we average across event time periods. The average coefficients are presented on Figure 8 (a). Apart from larger confidence bands due to the smaller number of observations, this figure is in every way comparable to the baseline estimates presented on Figure 1.

As last alternative sample is constructed by using a new measure of a patent's quality in the propensity scores (based on which we weight acquired and non-acquired patents). Instead of a simple average of all 7 patent quality's indicators, and in order to capture as much information about a patent's quality as possible, we use the principal components that maximise the variation in our data ($pc1$ and $pc2$, as described

Figure 6: Big Tech citations to acquired patents relative to the time of acquisition

(a) excluding outliers (b) reducing the study period

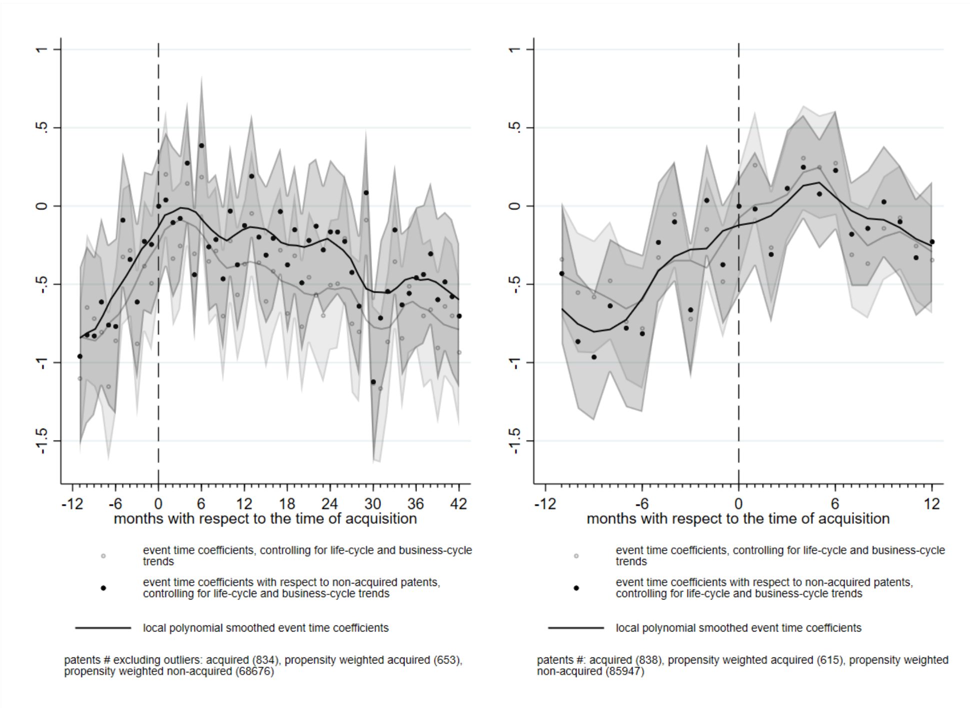

Notes: The graphs show the event time coefficients: θ^1 from model 2 (grey dots) and α^1 from model 3 (black dots). The event time coefficients on the left have been estimated on a sample from which outliers (i.e. patents associated with citations values whose expected frequency <0.5 assuming that they are NB-distributed) have been excluded. On the right, we reduced the study period from 4.5 to 2 years. The shaded bands represent 95% confidence intervals.

in Appendix III). As can be seen from Figure 8 (b), and compared to our baseline estimates, the event time coefficients are almost unaffected.

Figure 7: Propensity score distribution by acquisition status (Density)

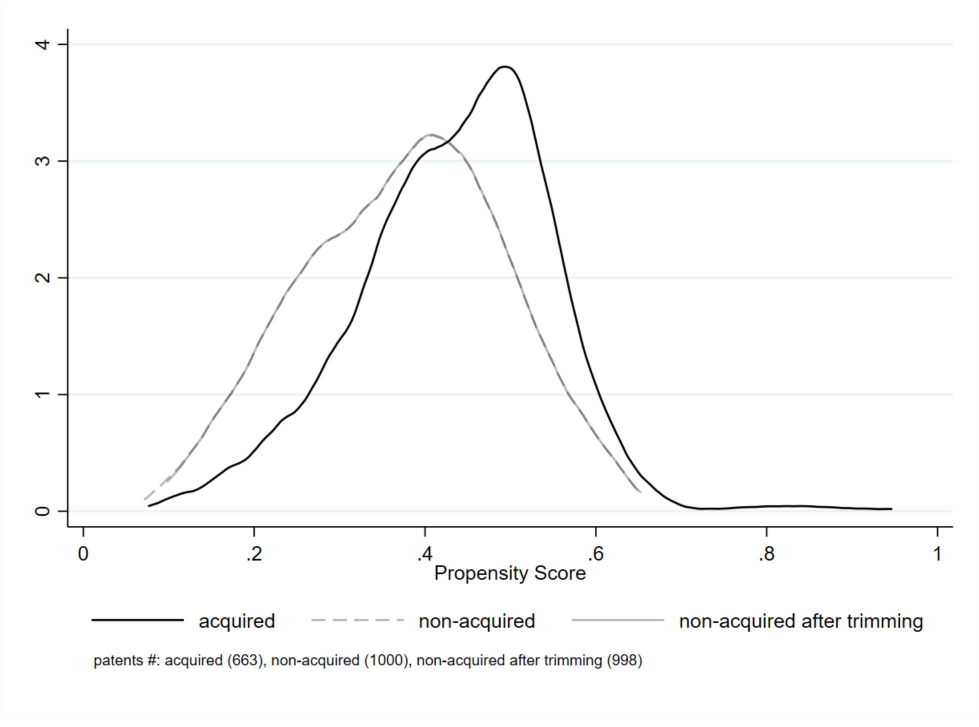

Notes: This figure shows the distribution of the propensity scores for the full sample of acquired and for 1,000 observations randomly drawn from the sample of non-acquired patents.

Figure 8: Big Tech citations to acquired patents relative to the time of acquisition

(a) subsampled non-acquired patents (b) principal component as quality measure

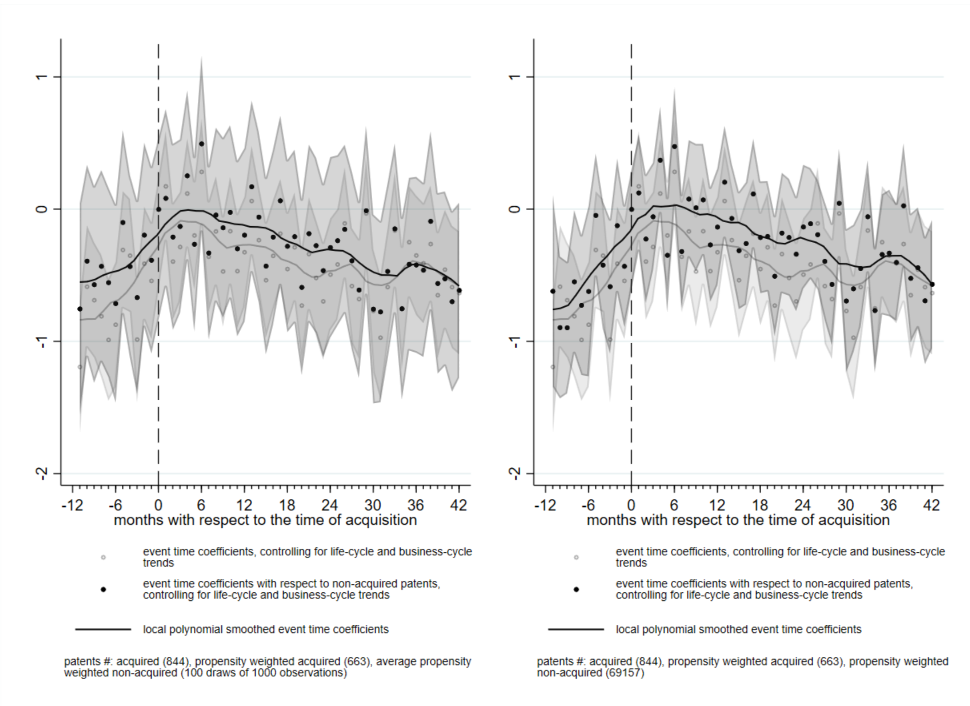

Notes: The graphs show the event time coefficients: θ^1 from model 2 (grey dots), and α^1 from model 3 (black dots). These coefficients are estimated on a balanced sample of patents in a 4.5 year-window around (simulated) acquisition. The shaded bands represent 95% confidence intervals.

Appendix VIII

Alternative model specifications

Figure 9: Big Tech citations to acquired patents relative to the time of acquisition

(a) acquisition announcement as event time (b) dispersion parameter constrained to zero

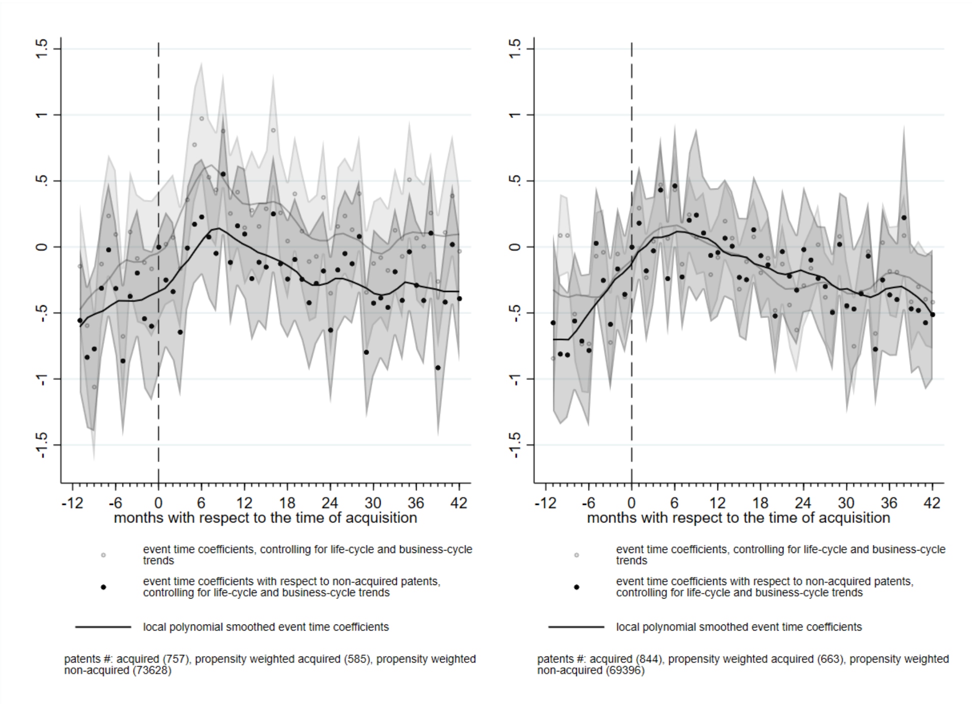

Notes: The graphs show the event time coefficients: θ^1 from model 2 (grey dots), and α^1 from model 3 (black dots). These coefficients are estimated on a balanced sample of patents in a 4.5 year-window around (simulated) acquisition. The shaded bands represent 95% confidence intervals.