

Big-tech M&A, Venture Capital and Innovation

Klaus Gugler, Florian Szücs & Ulrich Wohak*

First draft: May 2022

This draft: September 23, 2022

Disclaimer: This project is work in progress. Do not circulate!

Abstract We evaluate the impact of big-tech acquisitions on the incentives for investment and innovation. To this end, we collect data on several hundred big-tech acquisitions and compare venture capital investment in acquisition markets to a matched control group. Similarly, we study the evolution of patenting in technology classes with big-tech acquisitions. While the impact of a big-tech acquisition on investment is unequivocally negative, the effect on innovation depends on time dynamics and the identity of the acquirer.

*Klaus Gugler, Wirtschaftsuniversität Wien, klaus.gugler@wu.ac.at; Florian Szücs, Wirtschaftsuniversität Wien, florian.szuecs@wu.ac.at; Ulrich Wohak, Wirtschaftsuniversität Wien, ulrich.wohak@wu.ac.at.

1. Introduction

Acquisitions of young, start-up firms by large incumbents in technology markets have increasingly drawn the attention of policy makers in recent years. Five of the biggest tech companies - Google (now Alphabet), Amazon, Facebook (now Meta), Apple and Microsoft, collectively referred to as GAFAM - have jointly acquired more than 1,500 (???) companies since their foundation. Very few of these acquisitions have been scrutinized by competition authorities, as most transactions remain below the size thresholds specified in the Hart-Scott-Rodino (HSR) Act.¹ Only a handful of these acquisitions have been remedied or blocked.

While the combined market valuation of these five platforms is more than \$5 trillion (more than a third of the value of the S&P 100), they predominantly buy small nascent or potential competitors. This development has given rise to severe economic but also political and societal concerns about the functioning of competitive markets and the accumulation of assets and power. An important aspect of the vivid M&A activity of GAFAM firms is its effects on the innovative potential of economies. On the one hand, there are concerns that these acquisitions could lead to a decline in innovation, particularly in digital markets. "Killer acquisitions" (in order to shut down a future competitor, see Cunningham et al. (2021)), creation of "kill zones" (discouraging other firms from investing or innovating in specific areas, see Kamepalli et al. (2020) or conventional raising rival's cost strategies (e.g. by preventing other firms from access to complementary assets, see Salop and Scheffman (1983)) may follow these acquisitions and leverage market power to new markets and curb future innovative activity. This is the overwhelming concern in most competition reports,² the predominant finding in the empirical literature³ as well as the theoretical literature.⁴

On the other hand, there is a predominantly theoretical literature that also sees positive aspects of big-tech acquisitions.⁵ For example, Rasmussen (1988) conjectures that one of the reasons innovative firms are founded and funded in the first place is the prospect of being acquired by a large incumbent firm as the most important exit strategy ("entry for buyout"). Cabral (2020) stresses incomplete markets for technology transfer, and acquisitions may be a

¹Therefore, the Federal Trade Commission (FTC) issued Special Orders to these five companies requiring them to provide information about prior acquisitions not reported to the antitrust agencies under the HSR in the course of the large congressional investigation. See "Non-HSR Reported Acquisitions by Select Technology Platforms, 2010-2019: An FTC Study"

²See "Investigation of Competition in Digital Markets: Majority Staff Report and Recommendations, Subcommittee on Antitrust, Commercial and Administrative Law of the Committee on Judiciary", "Final Report of the Stigler Committee on Digital Platforms (September 2019)", "Unlocking Digital Competition: Report of the Digital Competition Expert Panel" (Furman report), "Ex-post Assessment of Merger Control Decisions in Digital Markets" (2019 LEAR report), "Australian Competition & Consumer Commission's Digital Platforms Inquiry" (July, 2019), "The French Competition Authority's Opinion on the Online Advertising Sector" (March, 2018), "A report to the European Commission, Competition Policy for the Digital Era", (Jacques Cremer, Yves-Alexandre de Montjoye and Heike Schweitzer, 2019).

³See e.g. Kamepalli et al. (2020) (negative effects on venture capital, creation of "kill zones"), Gautier and Lamesch (2021) (substantial portion of acquired products and services discontinued), Affeldt and Kesler (2021) (half of acquired apps in Google Play Store shut down), and Koski et al. (2020) (reduction of entry rates and venture capital funding in target's product market).

⁴See e.g. Motta and Peitz (2021) providing an overview over the mostly harmful competitive effects of such acquisitions.

⁵See also a recent empirical paper by Doan and Mariuzzo (2022) finding positive effects of leading firm acquisitions on patenting in the cloud computing market.

means for incumbents to source complementary assets. Their high willingness to pay for such assets creates innovation incentives for entrants. Letina et al. (2021) find in their model that a complete prohibition of large tech acquisitions would result in lower innovation efforts.

The above literature, however, suffers from a number of drawbacks which do not yet allow general conclusions on the effects of GAFAM mergers on innovation incentives. First, there is no comprehensive study of GAFAM acquisitions on innovation outcomes. For example, Kamepalli et al. (2020) analyze only nine large software acquisitions by Facebook or Google. Doan and Mariuzzo (2022) analyze only the cloud computing market, where only three out of the five GAFAMs are active (Amazon, Microsoft and Google). Conversely, in this paper we analyze the entirety of acquisitions on which data are available, by all five GAFAM firms.

Further, the outcome variables used in the extant empirical literature are not entirely convincing proxies for innovation. For example, venture capital (VC) funding, the number of products or the number of apps - while arguably related to innovation - have their deficiencies in measuring innovation input or output. VC funding may also include other, non-innovation outlays such as firm general purpose expenditures, while new products and apps are not necessarily innovative. We also analyze venture capital funding but complement it with an analysis of the patenting behavior of companies. While patent measures suffer from their own problems,⁶ they are closely related to innovation. We argue that patents are important measures of innovation in digital markets, as evidenced by the sheer size of the patent portfolios of GAFAM firms as well as the prevalence of patenting activity of their targets. Moreover, we stress the combination of our two measures of investment and innovation, VC funding and patenting activity. Acquisitions as well as their effects are a heterogeneous form of investment. One measure of innovation does not appear to fully capture this heterogeneity.

MORE

Thus, while the literature finds a preponderance of potentially negative effects of large tech acquisitions on the competitive process and in particular on innovation incentives, strong conclusions are not yet possible. We contribute to these questions by analyzing the effects of all GAFAM acquisitions on i) venture capital funding in the relevant markets and ii) patenting in the relevant technology classes. We interlink three sources of data and assemble a database containing (1) all acquisitions of GAFAM firms (SDC Platinum) from 1990 - 2020, (2) all firms with venture capital funding including GAFAM targets as well as other targets involved in other tech mergers (Crunchbase), and (3) all applied for and granted patents of all involved firms (PATSTAT).

The most tricky part in any study of the effects of mergers is to determine the proper counterfactual. Which innovative/VC activity would have been observed had the merger not taken place? This is a particularly difficult problem with tech mergers involving very young, small and nascent firms, where established counterfactuals simply do not exist. Fortunately, our

⁶Not all innovations are patented, and not all patents are (valuable) innovations. Relatedly, patents suffer from skewness in value (some patents are very valuable, most are not). There may be strategic patenting, e.g. for foreclosure reasons, "patent thickets", etc. See however (Lanjouw and Schankerman, 2004) for ways to increase the accurateness of patents as measure for innovation.

research question allows us to circumvent this problem in so far as we are not interested in the effects of mergers on innovation of a particular, specific firm (e.g. the target or the combined firm), but on their effects on a group of firms namely the "industry" or "space". Thus, we ask what are the effects of mergers on the number and citations of patents in the (8-digit) IPC class the merger takes place or the number and value of venture capital deals in the "space" the target is part of. Put differently, our treated units are not the specific targets or the combined firms, but the "industry" or "space" around the targets. Conceptually, a treated industry or space are the set of companies that are affected ("treated") by the acquisition, since these companies are reasonably similar to the target firm.⁷

Our approach alleviates two problems. First, as indicated, looking at the counterfactual for a treated industry or space of companies is – arguably – less controversial than determining the counterfactual of an acquisition for a particular firm (e.g. the combined firm). The prospective or counterfactual development of "spaces" of companies is much less heterogeneous, and much more predictable, than the counterfactual development of a particular firm. Second, and more importantly, we want to determine whether and which acquisitions added value to society, e.g. because they increased overall patenting and/or venture capital funding, and whether and which acquisitions destroyed value, e.g. because they have been "killer acquisitions" or created "kill zones", reducing overall innovation in the economy. To determine these competitive effects on *society*, we need to know the effects of these acquisitions on the whole industry or space of companies.⁸

Analyzing all acquisitions and two outcome variables related to investment and innovation allows us to achieve several objectives. First, we can draw a more nuanced picture of the innovative effects of GAFAM acquisitions than so far possible. Not all acquisitions are killer acquisitions nor do all acquisitions create kill zones. However, some do, and we try to describe and define those acquisitions, which do. If an acquisition reduces both, venture capital funding and patents in an industry, the negative effects of this acquisition are likely to outweigh any positive effects. Second, there is a lot of heterogeneity across acquisitions. Some acquisitions are unambiguously "bad" (if they reduce both metrics), however some are in between affecting one metric positively and the other metric negatively, and some acquisitions are even unambiguously "good" increasing both metrics in the industry. We characterize these types of acquisitions. Finally, we find heterogeneous effects across GAFAM acquirers. Our large sample allows us to draw conclusions on this dimension, too.

Our main results indicate that while the impact of a big-tech acquisition on VC funding is unequivocally negative, the effect on innovation (patents) depends on time dynamics and the identity of the acquirer.

The rest of this paper is organized as follows. Section 3 describes our dataset, while our

⁷We define "space" similarly to Kamepalli et al. (2020) capturing this reasonable degree of similarity.

⁸For example, it may be that an acquisition increases the number of patents for the combined firm, but lowers it for the whole industry. This would be detrimental to total welfare.

empirical strategy is detailed in section 4. Results of the empirical analysis are discussed in section 5 and section 6 concludes.

2. Theoretical predictions

Any theory of the effects of big tech mergers must incorporate the characteristics of the firms and industries involved. The standard set up involves a dominant digital platform (the "incumbent") acquiring a potential or nascent competitor (the "start-up"). The start-up is often small but quickly growing in adjacent markets. Crucially, if the start-up's project is successful it can become a substitute of the incumbent's product or service, i.e. it becomes an actual competitor in the future. Characteristics of digital industries include network effects, multi-sidedness, free provision of services, and the importance of data. Motta and Peitz (2020) provide a good starting point which effects may be expected depending on this set up. They stress the importance of the likely counterfactual, and show that whenever the start-up has the ability to pursue its project, the merger will be anti-competitive. To put it differently, the acquisition can only be pro-competitive if the potential competitor is unable to pursue the project absent the merger and if the incumbent has an incentive to pursue, rather than shelving, the project.

Several aspects of merger effects follow from this simple intuition. First, if the "Arrow replacement effect" does not hold (then duopoly profits are larger than monopoly profits), then there may exist projects that an entrant would not carry out, while an incumbent would. Motta and Peitz (2020) state that the Arrow replacement effect holds for several standard oligopoly models, but may not hold for general quality-enhancing innovations. In mergers involving targets that are active in adjacent markets developing complementary products (thereby accumulating a customer base, data or technology) which may however become substitutes in the future, the incumbent has an incentive to acquire this target. This results in early elimination of potential competitive threats.⁹

Under the assumption that monopoly profits are higher than the sum of duopoly profits (Arrow replacement effect holds), at equilibrium the incumbent will always bid more than another bidder and win the bidding competition. Thus, a high willingness to pay on the side of the incumbent may indicate that the incumbent may want to protect its incumbency rents.¹⁰

Motta and Peitz (2020) in an extension also endogenize the resources of the start up and allow for exclusionary conduct by the incumbent. Exclusionary practices (e.g. refusal to supply, degradation of interoperability, tying/bundling or imitation of the entrant's products) may be

⁹Cunningham et al. (2018) find that when incumbents acquire firms which have been developing competing drugs, they are more likely to abandon such projects if they possess more market power. Thus, the more competition there is among incumbents, the less likely it is that killer acquisitions take place.

¹⁰The acquisitions of Instagram and Whatsapp by Facebook and of Waze by Google to prevent competition in the social network apps market and with Google Maps, respectively, are often cited that they follow this pattern. A policy of preventing incumbent firms from bidding must however trade off the pro-competitive effects against possible negative innovation effects if expected payoffs of the innovator are reduced, see Fumagalli et al. (2020).

costly for the incumbent in the short run but may increase future profits.¹¹ Exclusionary conduct and acquisition may be complementary, since an exclusionary strategy reduces the acquisition price. Thus, the existence of a "kill zone" (according to which new firms/investors would stay away from the core market of large digital platforms) can be explained by the threat the incumbent may engage in exclusionary practices. Kamepalli, Rajan, and Zingales (2020) were the first to show that acquisitions by the incumbent may lower payoff prospects of new entrants and thus discourage them from investing. Choi and Jeon (2020) analyze bundling strategies and find that bundling can lead to a leveraging of monopoly power from the incumbent into the target market, creating a "kill zone" as more efficient competitors would decide not to enter the acquisition market.

Another theory of harm is related to the enhanced possibilities of collecting data after an acquisition. The idea is that with more consumer data the match quality between advertiser and consumer is improved. By increasing its quality after merger the incumbent market share is increased. This may lead competitors to provide less quality leading to a loss of market share. Thus, while the merger may be consumer welfare increasing in the short run, it may be harmful to consumers (and innovation) in the long term, as firms may be induced to exit.

Authors have not only looked at the quantitative innovation incentive effects of mergers, but also whether the prospect of being acquired changes the type of innovation undertaken. Bryan and Hovenkamp (2020) conjecture that the leader will always acquire the start-up to prevent the laggard from catching up technologically. Start-ups will then bias their R&D investment towards improving the incumbent's technology rather than towards technology helping a/the laggard to catch up.

Cabral (2020) somewhat in contrast to Motta and Peitz (2020) highlight the importance of technology transfer through acquisitions in the presence of imperfect knowledge markets. Digital industries are characterised by high uncertainty about where the next competitive threat comes from, which lowers the preemption motive for acquisitions. Acquisitions allow the transfer of complementary technology to the incumbent. If this technology is worth more for the incumbent, the higher acquisition price generates innovation incentives for entrants in the first place (see also Letina, Schmutzler and Seibel (2020)).

Of course, all theories of harm suggesting decreased innovation incentives must be traded off against any possible efficiency effects arising from the merger. These may include innovation complementarities, internalization of innovation externalities, positive network effects as well as other efficiency defences. Thus, generally likely effects of big tech mergers on innovative activity are ambiguous and a matter of empirical evidence.

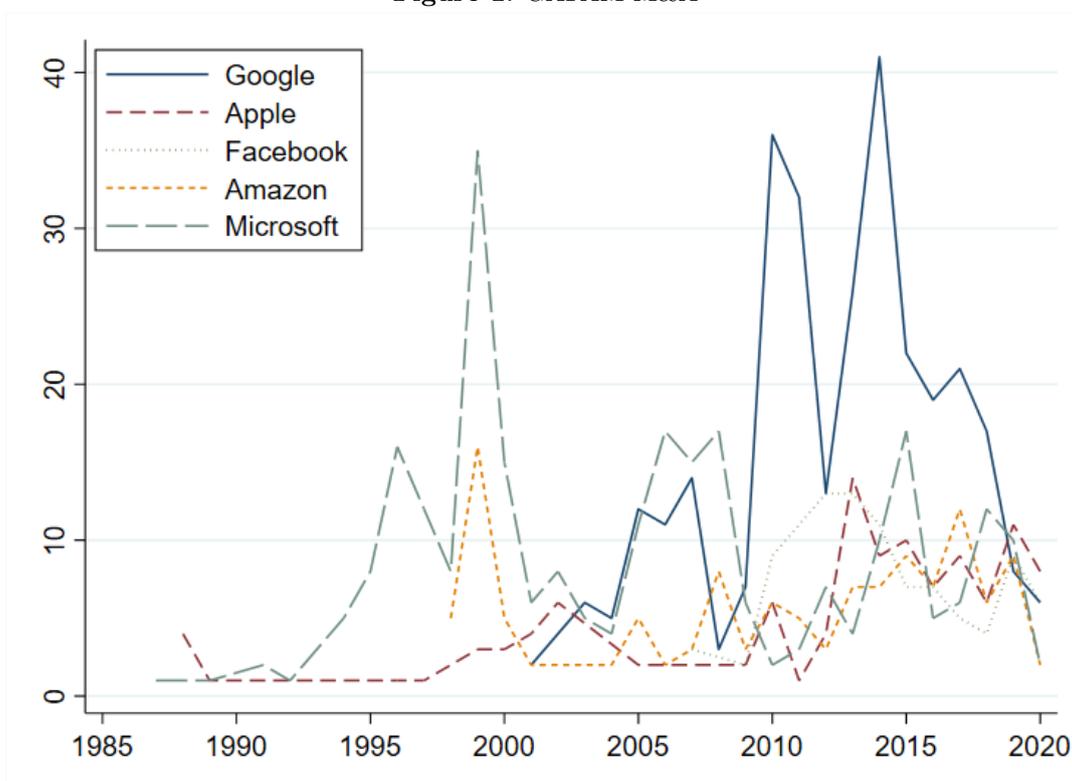
¹¹Exclusion may be all the more likely if there is no trade off between short and long run profits as possibly with some imitation strategies, see Motta and Shelegia (2020).

3. Data

3.1. Mergers and acquisitions

We use Thomson Reuters SDC Platinum to obtain a full list of all mergers and acquisitions associated with Google, Apple, Facebook, Amazon and Microsoft. We manually clean the results and remove entries that do not correspond to proper M&A (e.g. repurchases) and in addition use financial news websites to corroborate all deals. We discard transactions the status of which is 'pending' or 'unknown' and only retain 'completed' deals. We obtain a list of 912 M&A, 301 by Google, 116 by Apple, 100 by Facebook, 124 by Amazon and 271 by Microsoft. While the acquisitions range from 1987 to 2020, only 7% occur before 1999 and more than half occur in the 2010-2019 period (see figure 1). These M&A constitute the 'treatment' in our setting and we will study their impact on investment and innovation outcomes.

Figure 1: GAFAM M&A



3.2. Investment

To investigate the potential effects of GAFAM acquisitions on investment, we collect data on business activities of start-ups and innovative firms as well as venture capital financing activity from Crunchbase (<https://www.crunchbase.com>, accessed on 9 March 2021), a commercial database that collates information on firms with a particular focus on start-ups. We obtain data for 1.3 million firms on their identity, business description, total number of employees, total

venture capital investment as well as a list of industry classifications assigned by Crunchbase.¹² We use the firms' business description to identify firms that are similar to acquisition targets of GAFAMs.

Crunchbase also collects information on financing rounds, covering funding stage (Angel, Seed, Private Equity, etc.), identity of both recipient and investor, the amount invested as well as the date of the financing round. Together, this allows us to construct around each acquiree a "space" of similar firms and link early-stage VC investment to each space.

We are able to match 670 of the 912 GAFAM M&As in the Crunchbase data. For the resulting set of transactions, only 61 (9.11%) occurred before the year 2000. The vast majority of matched transactions took place between 2000 and 2020. 189 (28.21%) occurred between 2000 and 2010, while 420 (62.69%) occurred between 2010 and 2020. For each matched transaction, we obtain for all firms operating in the same space as the target (see Section 4.1) the year and amount of early-stage venture capital investment¹³. We drop acquisitions as well as investment deals that take place before 1990 since records in the Crunchbase data are incomplete for this period. We then aggregate this firm-specific early-stage investment as well as the number of early-stage VC deals in a 5-year window around each merger to yearly sums for each target, allowing us to track VC investment and deals for firms operating in the same space as GAFAM targets over time. We construct in a similar fashion counterfactual early-stage VC investment (see Section 4.1) resulting in a panel of 13,118 space-year observations of early-stage VC investment and deals.

3.3. Innovation

We use data from PATSTAT to track firms innovation activities. PATSTAT is maintained by the European Patent Office and contains data on worldwide patent applications and indicators.

We start by generating a balanced panel dataset at the IPC8-level. The International Patent Classification (IPC) sorts patents into fine-grained technology classes of up to 8-digits. We retain all patents granted in the 1980-2020 (the last year for which patent data are available) period in a total of almost 75,000 IPC classes, yielding more than 3 million observations. Thus, the unit of observation is an IPC class in a year and the resulting panel tracks the total amount of patents and the number of citations received in a 5-year window CITE FOR THIS, as well as some technology-class specific covariates, such as the number of inventors active, their share of total patents and the growth rate of patents, used for matching below.

Next, we match acquirers and targets to their respective patents. The large technology companies in the data account more than 400,000 patents in the sample period (Google 86k, Apple 112k, Facebook 22k, Amazon 40k and Microsoft 142k) and have patents in a total of 10,139 IPC8 technology classes.

¹²For a complete list of Crunchbase industries see <https://support.crunchbase.com/hc/en-us/articles/360043146954-What-Industries-are-included-in-Crunchbase->

¹³We classify venture capital investment as early-stage investment if the funding stage was "Angel", "Pre-seed", "Seed", "Series A" or "Series B". For a full list and description of funding types see: <https://support.crunchbase.com/hc/en-us/articles/115010458467-Glossary-of-Funding-Types>

Out of the 912 GAFAM M&A in the data, we are able to link 355 target firms to PATSTAT, i.e. 39% of targets have at least one, granted patent. The fact that this share is much higher than the population average indicates, that these are technology-focused acquisitions (for example, only around 5% of firms listed in the ORBIS database have any patents). Most acquisitions linked to patents are undertaken by Google and Microsoft (118 and 120 respectively), followed by Amazon and Apple (45 and 46 respectively), while there are only 26 deals by Facebook. In sum, the target firms have been granted a total of almost 90,000 patents in 5,665 different IPC8 technology classes.

We combine the IPC8-panel with the information on which big-tech company acquired targets. Thus the resulting dataset contains almost 75,000 IPC classes in the 1980 - 2020 period and indicates in which IPC classes and years big-tech acquisitions were made. As the majority of IPC classes were not affected by GAFAM deals, they serve as a donor pool for matching a control group below.

4. Empirical strategy

Here we describe how the data discussed in the previous section are set up for analysis.

4.1. Constructing investment markets

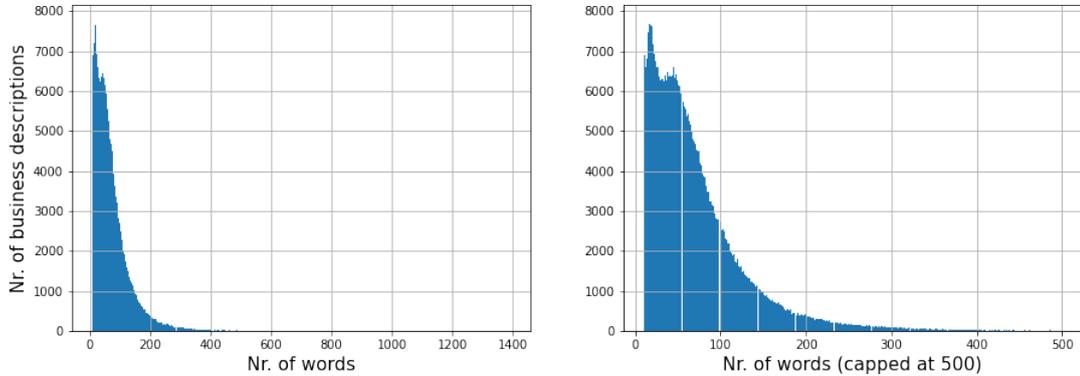
We characterize the space of "treated" firms as the set of firms that are sufficiently similar, and hence exert a competitive constraint to, GAFAM targets. Conventional economic market delineation is based on demand side (e.g. SSNIP) and/or supply-side substitution approaches using data on prices, quantities and other observable market outcomes. The challenge in applying this approach to nascent technology markets is the fact that such data are not available. In their absence, we leverage text data on firms' business descriptions. Crunchbase has data on 1,2 million firms of which 647,509 have a business description. For this subsample, the longest business description consists of 1,392 words while the shortest has 1 word. The median business description is 55 words long¹⁴. In the subsequent analysis we only consider firms with a business description of at least 10 words to ensure that descriptions are meaningful. This leaves us with 624,315 firms' business descriptions for further analysis.

We apply natural language processing (NLP) techniques to assess the similarity of firms' business activities to identify firms that are likely to operate in the same "space". The main challenge in the analysis of text data is the removal of noise associated with *polysemy* (multiple meanings for one word) and *synonymy* (multiple words for one meaning, e.g. "work", "occupation", and "profession"). To capture this underlying semantic structure we employ Latent Semantic Analysis (LSA) (Deerwester et al., 1990) to the business description of firms listed on Crunchbase. LSA is a natural-language processing method that retrieves semantic similarities between words

¹⁴Crunchbase provides both a "Description" and a "Short description". There are cases where, contrary to the name of these columns, "Short Description" contains more words than "Description". In such cases, we use the column with more words.

if they are used in similar contexts and hence allows us to assess the similarity of business description on the basis of their latent semantic content.

Figure 2: Distribution of business descriptions over number of words



Before applying LSA, we exclude stopwords¹⁵ and apply a "lemmatizer", a NLP tool that performs a morphological analysis on each word and reduces it to its lemma¹⁶¹⁷. From this corpus we construct a $D \times W$ document-term matrix A , where D is the set of documents and W is set of distinct words in the corpus. Each element $c_{i,j} \in A$ corresponds to the number of times word $w \in W$ occurs in document $d \in D$. We re-weight the entries in A by their term-frequency inverse-document frequency, such that words that appear frequently across documents are assigned a low weight whereas words that appear frequently in a single document are assigned a higher weight. This re-weighted matrix A is then decomposed using singular-value decomposition resulting in the best rank- C approximation $\mathbf{B}_C = \mathbf{U}_C \mathbf{\Sigma}_C \mathbf{M}_C^T$. Each document is represented by a $1 \times C$ vector v_i in the document-component matrix $\mathbf{U}_C \mathbf{\Sigma}_C$. Finally, we assess the similarity between documents using cosine-similarity between document vectors¹⁸.

The "space" around each target is a caliper that is defined by the cosine similarity of business descriptions between a GAFAM-target and other firms. For each GAFAM-target, firms with a cosine similarity greater than 0.3 are included in the space. There is no clear rule for choosing this particular cutoff, but visual inspection of the similarity in business description suggests that 0.3 is a sensible choice. We provide robustness results with respect to the cutoff in Appendix B. The median target-specific treated space is composed of 1002 firms, for which 415 early-stage investments are recorded.

We construct the counterfactual investment trajectory for each space using industry classifications from Crunchbase. Firms are categorized in industries such as "Augmented Reality", "Cloud Computing", "Machine Learning"¹⁹. The typical firm is assigned to more than one industry

¹⁵Stopwords are words that are considered to carry little semantic meaning, such as "and", "or", "this", "that", etc. In the context of our analysis, we extend this list by additional words that are likely to add little value, such as "LLC", "www", "GmbH", etc.

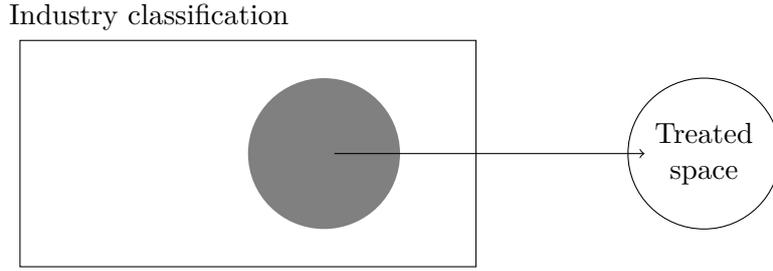
¹⁶For more information see: <https://nlp.stanford.edu/IR-book/html/htmledition/stemming-and-lemmatization-1.html>

¹⁷For example, the words "buy", "buys", "bought", "buying" are reduced by the lemmatizer to its lemma "buy".

¹⁸See Appendix XX for a detailed explanation.

¹⁹A full list of industry classifications can be found here: <https://support.crunchbase.com/hc/en-us/articles/360043146954-What-Industries-are-included-in-Crunchbase->

Figure 3: Illustration of treatment and control space



group. Therefore, we include in the counterfactual all firms that are active in at least one industry of the target. From this set of firms, we exclude all firms that belong to a target’s space. The median target-specific control space is composed of approximately 29,000 firms for which we have records on 35,563 early-stage investment rounds²⁰.

4.2. Matching technology classes

When looking at the different IPC classes in the dataset it becomes apparent, that (just like the underlying individual patents) patenting is very heterogeneous across classes. For example, while the median amount of patents (citations) in an IPC8 class and a year is 14 (49) on average, these numbers rise to 104 (1153) in IPC classes with GAFAM acquisitions. It seems plausible that tech-giants strategically select which targets and technologies they acquire (rather than random selection into treatment), and summary statistics reported in the first three columns of table 1 seem to confirm this. The first two columns contain the means of key metrics in IPC8 classes with GAFAM M&A and in the population, the third column contains the p -value of a t -test for equal means. As can be seen, classes with M&A tend to have more patents, more citations and more inventors active in them. We therefore need to account for this selection process in order to avoid bias.

Table 1: Biases before and after matching

	Before matching			After matching		
	Treated	Population	p -value	Treated	Control	p -value
Patents	584.42	91.17	0.00	371.58	365.21	0.72
5-year citations	2604.34	244.87	0.00	1706.07	1773.06	0.53
Inventors	8.45	6.72	0.00	8.23	7.99	0.00
C1	-4.34	-3.97	0.00	-4.40	-4.38	0.26
High citation IPC	0.59	0.09	0.00	0.51	0.53	0.11
Trend	0.38	0.43	0.00	0.33	0.34	0.17
Patent growth	0.15	0.37	0.00	0.12	0.14	0.24

Notes: The number of inventors and the share of patents belonging to the largest patentor in a class (C1) are in logs. 'High citation IPC' indicates IPCs in the 90th percentile of the distribution of citations. Patent growth is the measured from one period to the next, 'trend' is the class-specific average over time.

²⁰Note that a firm might receive VC funding more than one.

We do so by employing a propensity-score matching procedure, in which a binary indicator for whether an IPC8 class is affected by GAFAM deals or not is explained through a set of innovation-related indicators and fixed-effects. The predicted values of this model, the propensity scores, then represent the ex-ante probability that a technology class is affected by GAFAM acquisitions. We then pair treated and non-treated IPCs with similar propensity scores in order to obtain a balanced sample.

We use a logit model to regress an indicator for GAFAM acquisitions in an IPC8 class on potential determinants. An IPC8 class is treated if a GAFAM firm acquired a target with patents in that class. Thus, a single acquisition can affect multiple technology classes. The estimation results are reported in table 2. While patent count remains insignificant (likely due to its high correlation with citations, $\rho = 0.76$), the total number of received citations in an IPC8 class and a year increases the likelihood of a GAFAM acquisitions (both patents and citations were logged for this estimation to avoid tiny coefficients). The log number of inventors has a positive impact on M&A, the log C1 (defined as the share of patents belonging to the largest patentor in a specific class), has a negative impact. High citation IPC (defined as the 90th percentile of total citations) are more likely to experience M&A. The current and average patent growth rate (patent growth and trend) have a negative and a positive sign, respectively.

Table 2: Selection model

Patents	-0.049	(0.033)
5-year citations	0.253***	(0.024)
Inventors	1.263***	(0.030)
C1	-0.382***	(0.031)
High citation IPC	0.198***	(0.052)
Trend	2.435***	(0.056)
Patent growth	-0.220***	(0.022)
Observations	1269189	
Pseudo R^2	0.476	

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

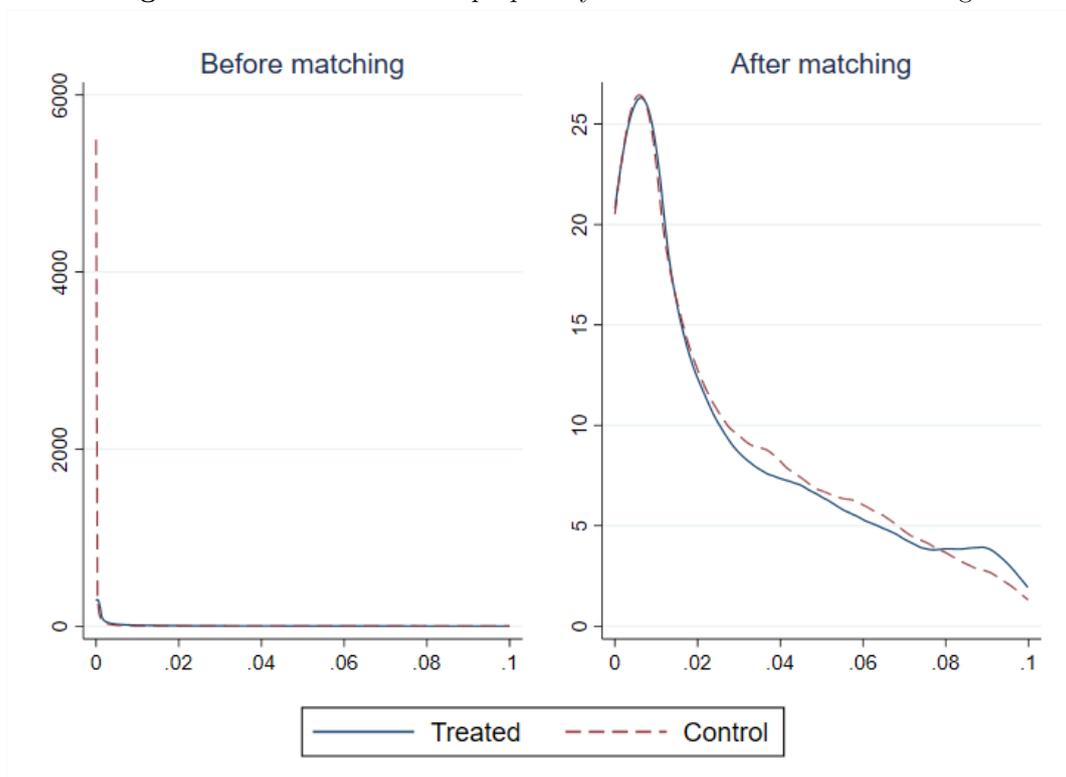
From this model, we obtain the propensity scores and proceed to match a control group. For each IPC8 class with GAFAM M&A, we indicate all other IPC8 classes at the same point in time, that have not been affected by M&A and have not already been assigned as a control unit (matching without replacement). We further require the potential controls to be in the same quartile of patent count and growth across IPC classes, to help with balancing. Among those candidate classes, we choose the one with the lowest (absolute) difference in propensity score as a control observation. Iterating this procedure across all IPC classes with GAFAM M&A yields 3,829 treated/control IPC8 pairs and a total of approximately a quarter million observations.

Figures 4 and 5 illustrate the distribution of propensity scores before and after matching, and the standardized biases²¹ in covariates before and after matching, respectively. While the

²¹Rubin (2001) proposes that the standardized bias (defined as $\frac{\bar{X}^1 - \bar{X}^0}{\frac{1}{2}(\sigma^1 + \sigma^0)}$, that is the difference in means of

distribution of propensity scores in the donor pool consists essentially of a mass point at (almost) zero (left side of figure 4), the two distributions look very similar after matching (right side of figure 4).

Figure 4: Distribution of the propensity score before and after matching



Similarly, while standardized biases are substantial prior to matching (figure 5), they mostly vanish in the matched sample. An exception to this is the number of inventors, which retains a standardized bias of 0.2 and is also significantly different between groups (see columns 4 - 6 of table 1). However, all other covariates and - most importantly - the dependent variable, 5-year citations, are balanced as can be seen from the t -tests in table 1. The remaining bias in the number of inventors is not too large (Rubin (2001) recommends post-matching biases to be <0.25) and is not likely to cause econometric concern.

We have thus created control groups for the affected investment markets and for the affected technology classes. We will use these counterfactuals to estimate, how these markets would have evolved in absence of GAFAM M&A.

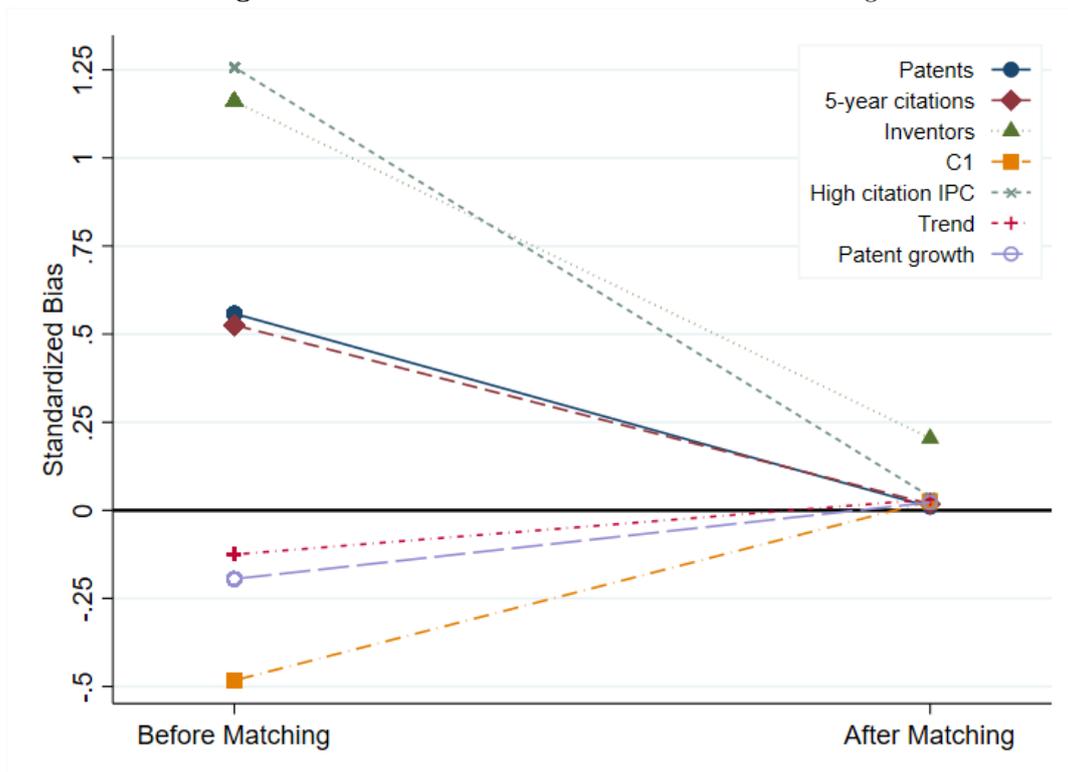
4.3. Estimation

The average treatment effect of big-tech acquisitions on investment and innovation behaviour is estimated in a DiD framework. The estimation equation is given by

$$\mu_{i,t} = \delta \text{post}_t + \gamma(\text{treated}_i \times \text{post}_t) + \iota_i + \tau_t + \varepsilon_{i,t},$$

treated and control groups divided by the mean of their standard deviations) of covariates should be below 0.25 after matching.

Figure 5: Biases in covariates before and after matching



where $\mu_{i,t}$ is the outcome (either log 5-year citations in an IPC8 class or log investment in a caliper) of unit i in year t , ν_i and τ_t are unit- and time-fixed-effects respectively. The error term is robust to heteroskedasticity and allowed to cluster at the IPC8-class level.

5. Results

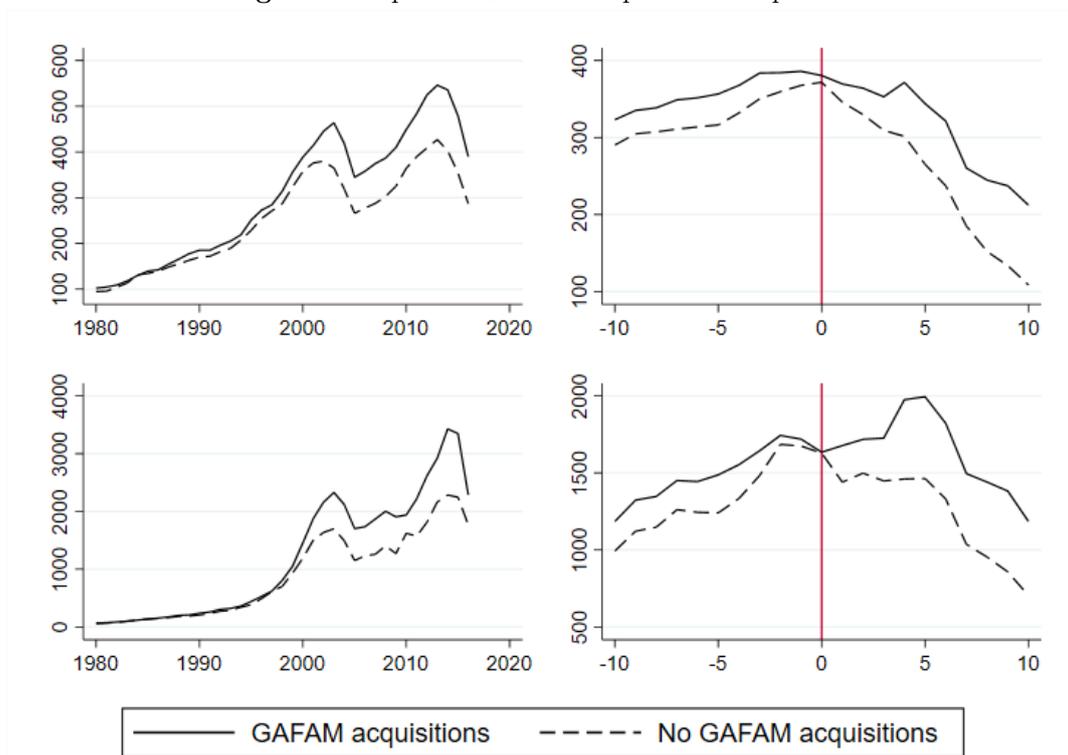
5.1. Main results

Figures 6 and 7 provide a general overview over patent counts, patent citations, VC deals and VC investment in both types of "spaces" of firms, treated and untreated. In the case of patents, treated spaces are those 8-digit IPC classes where at least one GAFAM firm has made an acquisition, while untreated spaces are those 8-digit IPC classes where there was no such acquisition but which are "similar" (according to our matching procedure) to treated 8-digit IPC classes. In the case of VC investments, treated spaces are similar calipers around the target of a GAFAM acquisition, while the untreated space is the general VC investment industry. The respectively left panels convey that all four metrics experience dramatic growth over our sample period from the beginning of the 1990ies until 2020 ???. This growth is fairly similar for our treated and untreated spaces of firms as defined above.

Judging from the respectively right panels and visual inspection, trends for treated and untreated groups are very similar for all four metrics, patent counts, patent citations, VC deals as well as VC investment. Formal tests corroborate the finding that patenting and investment

patterns of treated and untreated spaces display parallel trends. ???²². Thus, the condition for causal treatment effects that treatment and control firms behave similarly before acquisition is met. Visual inspection of the respectively right panels moreover already hint at the main results of this paper: While VC investment unambiguously goes down after acquisition in treated calipers relative to the control group, the situation may be different for patents.

Figure 6: Impact of GAFAM acquisitions on patents



Notes: The two upper panels show patent counts in absolute time (left) and relative to GAFAM acquisitions (right). The two lower panels report the same for patent citations.

Tables 1 and 2 shed light on the quantitative effects for patent citations (Table 1) and VC investment (Table 2).²³ All regressions control for a full set of fixed effects (treated and untreated calipers/IPC classes; and year fixed effects). While the effect of GAFAM acquisitions is indistinguishable from zero for patent citations, it is significant and economically large for VC investment. Treated calipers of firms lose 37,8% of their VC investment in the period after a GAFAM acquisition compared to average VC investment.²⁴

These main results mask however important heterogeneities across time and GAFAM firms. Columns 2 of the tables report treatment effect estimates of acquisitions undertaken before 2011, and from 2011-2020. There have been ??? GAFAM acquisitions before 2011 and ??? acquisitions from 2011-2020.²⁵ GAFAM acquisitions before 2011 significantly reduce both patent citations in treated 8-digit IPC classes (by 19.7%) as well as VC investment in treated calipers (by 53.9%).

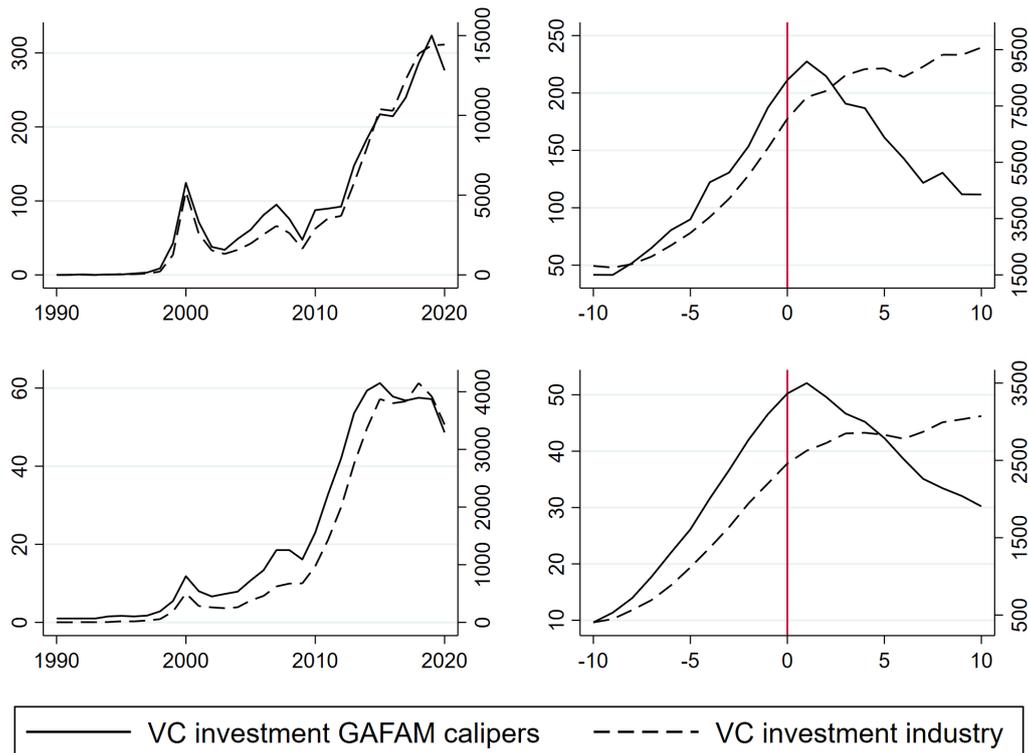
²²In regressions similar to Table ??? we regress patents/VC investment on yearly time dummies before acquisition for treated and untreated spaces of firms and find no significant differences using F-tests.

²³Results on patent counts and VC deals are comparable and available upon request.

²⁴The formula for converting natural log coefficient estimates into percent changes of odds is $(\text{Odds ratio} - 1) * 100$, where $\text{Odds ratio} = \exp(\text{betahat})$.

²⁵We take these two time periods as differentiating feature because the number of analyzable acquisitions before and after is reasonable allowing us to draw significant inferences. The trend results hold up in several other

Figure 7: Impact of GAFAM acquisitions on VC investment



Notes: The two upper panels show the amount of VC deals in absolute time (left) and relative to GAFAM acquisitions (right). The two lower panels report the same for the amount of money raised.

This contrasts to acquisitions consumed after 2010. Acquisitions by GAFAM firms undertaken between 2011 and 2020 increase patent citations in treated IPC classes by ???%. The same acquisitions reduce VC investment in treated calipers by ???%. Thus, while VC investment is still negatively affected by GAFAM acquisitions in post 2010 acquisitions, there is a clear difference in effects over time, in that effects become more benevolent to society over time.

5.2. The good, the bad and the ugly?

Before we can conclude that GAFAM acquisitions became "good" over time, we need to look closer into the heterogeneity of firms and deals. Table ??? presents the results on interaction terms of individual GAFAM firms and the time periods before 2011 and 2011-2020.

For Microsoft and Google acquisitions both metrics become more benevolent to society over time. While patent citation effects are very negative for both firms in the before 2011 period (-25 to -33% in treated 8-digit IPCs), they become both positive in the post 2010 period (+17.7% for Google and +97.8% for Microsoft). As for VC investment, both Microsoft and Google acquisitions display less negative effects over time. However, these types of acquisitions still crowd out between 18.4% (Google) and 31.3% (Microsoft) of VC investment in treated calipers.

At least one metric of the patent and VC effects of the acquisitions of the other three firms

subdivisions such as yearly or bi-annual dummies or dividing into three or four or five time periods. Thus, our inferences are robust to many other subdivisions of time periods.

Table 3: Patent citations and GAFAM M&A

	(1)	(2)	(3)
M&A	0.074*** (0.025)		
M&A_2010		-0.136*** (0.039)	
M&A_2020		0.303*** (0.028)	
Google			0.139*** (0.032)
Apple			-0.018 (0.051)
Facebook			0.565*** (0.069)
Amazon			0.113** (0.057)
Microsoft			-0.043 (0.042)
Observations	262527	262527	262527
R^2	0.678	0.679	0.679

Notes: Standard errors in parentheses, * p<0.1, ** p<0.05, *** p<0.01.

Table 4: Early-stage VC investment and GAFAM M&A

	(1)	(2)	(3)
M&A	-0.206*** (0.038)		
M&A_2010		-0.307*** (0.073)	
M&A_2015			
M&A_2020		-0.156*** (0.044)	
Google			-0.145** (0.068)
Apple			-0.283*** (0.102)
Facebook			-0.178* (0.093)
Amazon			-0.016 (0.120)
Microsoft			-0.386*** (0.072)
Observations	12371	12371	12371
R^2	0.934	0.934	0.934

Notes: Standard errors in parentheses, * p<0.1, ** p<0.05, *** p<0.01.

Table 5: VC and patent citations by firm and period

	Venture capital		Patent citations	
Google 2010	-0.322***	(0.119)	-0.267***	(0.061)
Google 2020	-0.064	(0.082)	0.200***	(0.035)
Apple 2010	-0.556**	(0.253)	0.206	(0.127)
Apple 2020	-0.206*	(0.108)	-0.107**	(0.048)
Facebook 2010	0.100	(0.145)	-0.105	(0.162)
Facebook 2020	-0.209**	(0.103)	0.561***	(0.078)
Amazon 2010	0.445*	(0.265)	0.111	(0.090)
Amazon 2020	-0.205*	(0.113)	0.061	(0.060)
Microsoft 2010	-0.529***	(0.105)	-0.396***	(0.056)
Microsoft 2020	-0.225**	(0.096)	0.737***	(0.070)
Observations	12371		262527	
R^2	0.935		0.680	

Notes: Standard errors in parentheses, * p<0.1, ** p<0.05, *** p<0.01.

Table 6: Acquisition-specific effects

	Acquisition-specific coefficients	
	Patents positiv	Patents negativ
Investment positiv	57	14
Investment negativ	89	37

display a deterioration. While Apple acquisitions decrease the negative VC effect (to (still) -27.1%), the patent citation crowding out in the treated 8-digit IPC class actually gets worse (to - 12.2%) in the second period (2011-2020). While patent citations are not crowded out before and after 2010 by Amazon acquisitions, VC investment starts to be crowded out in the 2011-2020 period (by -19.5%). While Facebook acquisitions increase patent citations (by 71%) in the second period, VC investors still flee the calipers where Facebook acquires a company (VC investment reduction of 37%).

Summarizing the heterogeneity results across individual GAFAM firms and time, we conclude that patent citations effects on treated 8-digit IPC classes become more benevolent to society with the exception of Apple acquisitions over time. VC investment effects become more benvolent for Google, Microsoft and Apple acquisitions, and deteriorate for Amazon acquisitions. However, all estimated VC coefficients are still negative in the post 2010 period for all GAFAM firms. Thus, while patents paint a heterogenous picture both across GAFAM firms and time, VC investors fear to lose their money when investing in similar to acquired firms.

6. Conclusion

References

- Affeldt, Pauline and Reinhold Kesler**, “Big Tech acquisitions—Towards empirical evidence,” *Journal of European Competition Law & Practice*, 2021, 12 (6), 471–478.
- Cunningham, Colleen, Florian Ederer, and Song Ma**, “Killer acquisitions,” *Journal of Political Economy*, 2021, 129 (3), 649–702.
- Deerwester, Scott, Susan T. Dumais, George W. Furnas, Thomas K. Landauer, and Richard Harshman**, “Indexing by latent semantic analysis,” *Journal of the American Society for Information Science*, 1990, 41 (6), 391–407.
- Gautier, Axel and Joe Lamesch**, “Mergers in the digital economy,” *Information Economics and Policy*, 2021, 54, 100890.
- Kamepalli, Sai Krishna, Raghuram Rajan, and Luigi Zingales**, “Kill zone,” Technical Report, National Bureau of Economic Research 2020.

- Koski, Heli, Otto Kässi, and Fabian Braesemann**, “Killers on the road of emerging start-ups—implications for market entry and venture capital financing,” Technical Report, ETLA Working Papers 2020.
- Lanjouw, Jean O and Mark Schankerman**, “Patent quality and research productivity: Measuring innovation with multiple indicators,” *The Economic Journal*, 2004, 114 (495), 441–465.
- Letina, Igor, Armin Schmutzler, and Regina Seibel**, “Killer acquisitions and beyond: policy effects on innovation strategies,” *University of Zurich, Department of Economics, Working Paper*, 2021, (358).
- Martin, Dian I and M Berry**, “Mathematical foundations behind latent semantic analysis,” *Handbook of Latent Semantic Analysis*, 2007, pp. 35–54.
- Motta, Massimo and Martin Peitz**, “Big tech mergers,” *Information Economics and Policy*, 2021, 54, 100868.
- Řehůřek, Radim and Petr Sojka**, “Software Framework for Topic Modelling with Large Corpora,” in “Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks” ELRA Valletta, Malta May 2010, pp. 45–50. <http://is.muni.cz/publication/884893/en>.
- Rubin, Donald B**, “Using propensity scores to help design observational studies: application to the tobacco litigation,” *Health Services and Outcomes Research Methodology*, 2001, 2 (3), 169–188.
- Salop, Steven C and David T Scheffman**, “Raising rivals’ costs,” *The American Economic Review*, 1983, 73 (2), 267–271.

A. Appendix: LSA

We use the information contained in companies' business description to identify companies that engage in similar business activities and are thus likely to compete against one another. We apply a method from the natural language processing literature: Latent Semantic Analysis (LSA)²⁶. To that end, we represent all business descriptions ("documents") in a document-term matrix \mathbf{A} where each row corresponds to a document $d \in D$ and each column corresponds to a word $w \in W$, where D and W correspond to the set of all documents ("corpus") and the complete vocabulary of the corpus, respectively. Each element in \mathbf{A} corresponds to the number of times word w occurs in document d :

$$\mathbf{A}_{D \times V} = \begin{pmatrix} c_{1,1} & \dots & c_{1,W} \\ \vdots & \ddots & \vdots \\ c_{D,1} & \dots & c_{D,W} \end{pmatrix}$$

We re-weight the individual elements in the document-term matrix \mathbf{A} by their term-frequency inverse-document frequency, as is standard in the NLP literature. This weighting assigns a low weight to words that appear frequently across documents and a high weight to words that appear frequently within only a few documents. In particular, each element in \mathbf{A} is re-weighted by $tfidf(c_{i,j}) = (1 + \log(c_{i,j})) * \left(\log\left(\frac{1+D}{1+d_w}\right) + 1\right)$ where d_w is the number of documents containing word w .

$$\mathbf{B}_{D \times V} = \begin{pmatrix} tfidf(c_{1,1}) & \dots & tfidf(c_{1,W}) \\ \vdots & \ddots & \vdots \\ tfidf(c_{D,1}) & \dots & tfidf(c_{D,W}) \end{pmatrix}$$

This re-weighted matrix \mathbf{B} is then decomposed using singular value decomposition (SVD). SVD transforms B of rank r into three matrices (Martin and Berry, 2007)

$$\mathbf{B}_{D \times W} = \mathbf{U}\mathbf{\Sigma}\mathbf{M}^T$$

where \mathbf{U} is a $D \times r$ orthogonal matrix, $\mathbf{\Sigma}$ is a $r \times r$ diagonal matrix and \mathbf{M}^T is an $r \times W$ orthogonal matrix. The final step in LSA is the elimination of $(r - C)$ rows and columns in $\mathbf{\Sigma}$ corresponding to the smallest eigenvalues of $\mathbf{\Sigma}$, where C is a scalar chosen by the researcher. There is no optimal scalar C , but it is recommended to choose $100 \leq C \leq 1000$ (Martin and Berry, 2007). In our application, $C = 500$. This truncation process results in the best rank- C approximation $\mathbf{B}_C = \mathbf{U}_C\mathbf{\Sigma}_C\mathbf{M}_C^T$ to the input matrix \mathbf{B} . Each document d_i is represented as a $1 \times C$ vector v_i in the document-component matrix matrix $\mathbf{U}_C\mathbf{\Sigma}_C$. We construct a measure of similarity between documents by calculating the cosine similarity between any two pairs of vectors v_i, v_j :

²⁶We implement LSA in Python using gensim (Řehůřek and Sojka, 2010)

$$S(v_i, v_j) = \frac{\sum_{c=1}^C (v_i \cdot v_j)}{\sqrt{\sum_{c=1}^C v_i} \sqrt{\sum_{c=1}^C v_j}}$$

where $S(v_i, v_j) \in [-1, 1]$ increases in the similarity of documents.

Figure 8: Impact of GAFAM acquisitions on patent citations

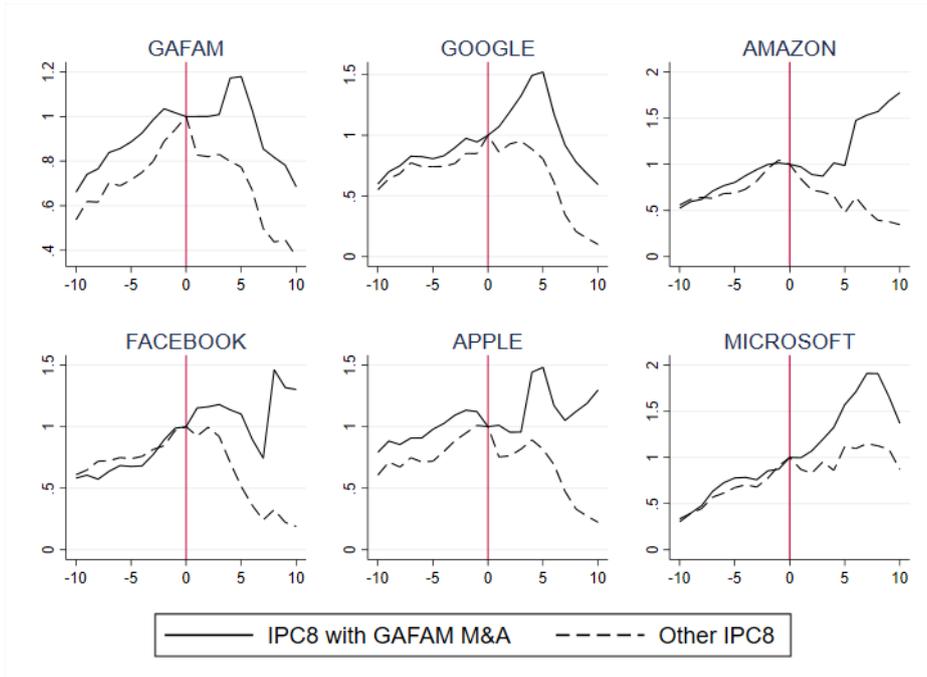
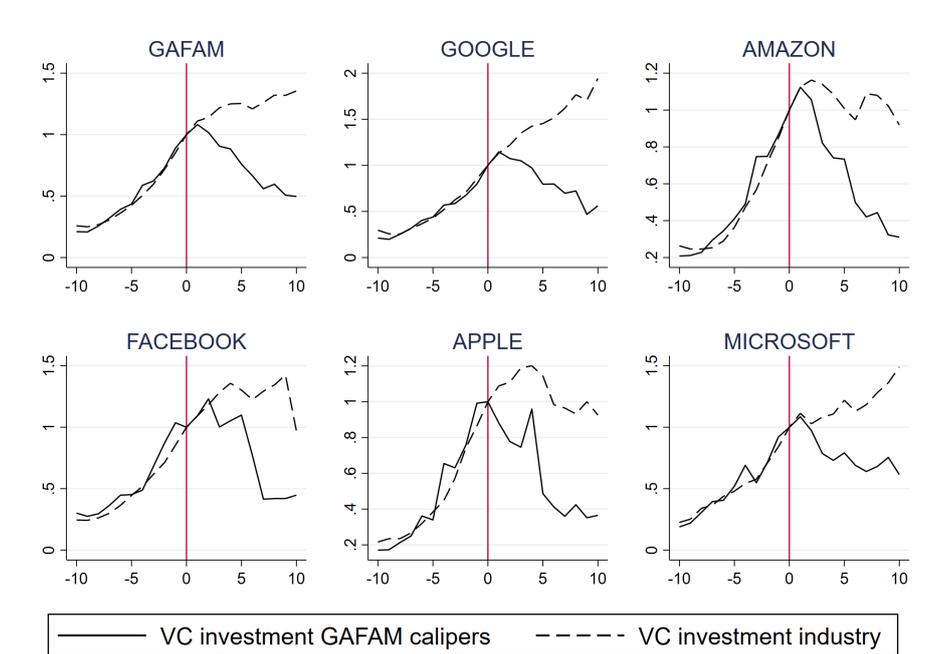


Figure 9: Impact of GAFAM acquisitions on VC investment



B. Appendix: Robustness cosine similarity

Industry	N	(+, +)	(+, -)	(-, +)	(-, -)
software	127	.32	.44	.06	.18
internetservices	65	.25	.58	.02	.15
hardware	54	.37	.37	.04	.22
informationtechnology	50	.28	.5	.02	.2
dataandanalytics	43	.42	.4	.05	.14
mediaandentertainment	42	.14	.64	0	.21
mobile	42	.24	.48	.1	.19
scienceandengineering	34	.44	.29	.12	.15
other	25	.28	.52	.04	.16
artificialintelligence	21	.57	.29	.1	.05
consumerelectronics	21	.38	.29	.1	.24
video	18	.22	.67	0	.11
apps	17	.24	.47	.06	.24
salesandmarketing	16	.25	.44	.06	.25
privacyandsecurity	15	.4	.47	0	.13
messagingandtelecommunications	14	.07	.64	0	.29
design	11	.27	.45	0	.27
commerceandshopping	10	.4	.3	.1	.2
professionalservices	10	.3	.5	0	.2
platforms	10	.2	.4	.1	.3
manufacturing	10	.1	.3	.1	.5
contentandpublishing	10	0	.8	0	.2

Notes: Crunchbase industry classifications with less than 10 GAFAM acquisition events are not displayed.

Table 7: Early-stage VC investment for different cosine similarity cutoffs

	s=0.3	s=0.4	s=0.5	s=0.6
M&A_2010	-0.307*** (0.073)	-0.643*** (0.106)	-0.731*** (0.146)	-1.083*** (0.304)
M&A_2020	-0.156*** (0.044)	-0.299*** (0.065)	-0.347*** (0.096)	-0.131 (0.173)
Observations	12371	9998	5496	2097
R^2	0.934	0.927	0.933	0.928

Notes: Standard errors in parentheses, * p<0.1, ** p<0.05, *** p<0.01.

Table 8: VC for different cosine similarity cutoffs by firm and period

	Venture capital			
	s=0.3	s=0.4	s=0.5	s=0.6
Google 2010	-0.322*** (0.119)	-0.648*** (0.157)	-0.606*** (0.199)	-1.054*** (0.405)
Google 2020	-0.064 (0.082)	-0.327*** (0.114)	-0.392** (0.159)	-0.519** (0.252)
Apple 2010	-0.556** (0.253)	-0.818** (0.335)	-0.525 (0.370)	-0.807* (0.441)
Apple 2020	-0.206* (0.108)	-0.302** (0.149)	0.134 (0.229)	0.541* (0.316)
Facebook 2010	0.100 (0.145)	0.016 (0.287)	-1.073*** (0.297)	-4.799*** (1.229)
Facebook 2020	-0.209** (0.103)	-0.134 (0.159)	-0.654*** (0.235)	-0.026 (0.539)
Amazon 2010	0.445* (0.265)	-0.100 (0.432)	-1.677*** (0.375)	-0.477 (0.731)
Amazon 2020	-0.205* (0.113)	-0.227 (0.155)	-0.267 (0.243)	0.186 (0.499)
Microsoft 2010	-0.529*** (0.105)	-0.900*** (0.182)	-0.707** (0.303)	-0.636 (0.486)
Microsoft 2020	-0.225** (0.096)	-0.468** (0.184)	-0.385* (0.204)	-0.027 (0.337)
Observations	12371	9998	5496	2097
R^2	0.935	0.927	0.933	0.930

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