The Aggregate Effects of Acquisitions on Innovation and Economic Growth*

Christian Fons-Rosen† Pau Roldan-Blanco‡ Tom Schmitz§

First draft: February 2021

This version: October 28, 2021

Click here to access the most recent version

Abstract

Large incumbent firms routinely acquire startups. The effect of these acquisitions on innovation and productivity growth is a priori unclear. On the one hand, acquisitions provide an incentive for startup creation, and a way to transfer ideas to potentially more efficient users. On the other hand, incumbents might "kill" some ideas of their targets, and acquisitions may create a less competitive environment with lower incentives for innovation. Our paper quantitatively assesses the net effect of these forces. To do so, we build an endogenous growth model with heterogeneous firms and acquisitions. We discipline the model by matching evidence from a rich micro dataset on startup acquisitions and patenting. Our calibrated model implies that acquisitions do raise the startup rate, but lower incumbents' own innovation as well as the percentage of implemented startup ideas. The negative forces are slightly stronger. Thus, a startup acquisition ban would increase growth by 0.03 percentage points per year.

Keywords: Acquisitions, Innovation, Productivity growth, Firm dynamics.

JEL Classification: O30, O41, E22.

*We thank Laurent Cavenaile, Florian Ederer, Chiara Fumagalli, Chad Jones, Enrico Moretti, seminar participants at Banco de España, Bocconi, UC Berkeley, UC Santa Cruz and UPF-CREi, and conference participants at the 2021 SED Annual Meeting and EEA-ESEM Virtual 2021 for helpful comments. The views expressed in this paper are those of the authors and do not necessarily coincide with those of the Banco de España or the Eurosystem.

†University of California in Merced, CEPR and MPI Munich, 5200 North Lake Rd. Merced, CA 95343, United States. Contact: cfons-rosen@ucmerced.edu.

‡Bank of Spain, Calle Alcalá 48, 28014 Madrid, Spain. Contact: pau.roldan@bde.es.

§Bocconi University, CEPR and IGIER, Via Roentgen 1, 20136 Milan, Italy. Contact: tom.schmitz@unibocconi.it.

1 Introduction

Large incumbent firms routinely acquire startups. For instance, the Tech giants Amazon, Apple, Facebook, Google and Microsoft have acquired at least 770 firms since their foundation.[1] Moreover, even though Tech acquisitions have recently captured the headlines, startup acquisitions are common in other industries as well.

Startup acquisitions are, however, viewed with increasing scepticism by regulators. In the United States, the Federal Trade Commission (FTC) recently announced an inquiry into several high-profile cases, and subsequently filed lawsuits against Facebook and Google.[2] While these inquiries traditionally focus on competition and prices, regulators have recently also started to worry about the effects of startup acquisitions on innovation. However, these effects are not obvious a priori. On the one hand, acquirers may choose to sideline startup innovations that threaten their existing business, and this could slow down productivity growth. On the other hand, the prospect of being acquired may stimulate startup creation, and actual acquisitions could improve the allocation of ideas between firms, which might accelerate productivity growth. Finally, acquisitions lower competition, which has ambiguous effects on innovation and growth.

In this paper, we aim to assess the relative strength of these forces. To do so, we develop a Schumpeterian growth model with heterogeneous firms and acquisitions. We discipline the model by matching empirical evidence from a new panel data set on acquisitions and patenting in the United States. Our calibrated model implies that the negative forces slightly dominate, so that reducing the frequency of acquisitions (e.g., through stricter antitrust enforcement) would lead to a modest increase in aggregate productivity growth.[3]

Our model builds on the Schumpeterian endogenous growth framework. Each incumbent firm produces a differentiated product, and seeks to innovate in order to increase its productivity. A large mass of non-producing startups, in turn, seeks to innovate in order to displace incumbents and enter the market. We introduce two novel elements into this setting. First, we distinguish between invention and implementation. That is, all firms first need to invest into invention (or research) in order to come up with an idea. After obtaining an idea, they then need to invest additional resources to implement it. Second, we allow for startup acquisitions: when a startup has developed an idea that could displace

[1] See https://www.cbinsights.com/research/tech-giants-billion-dollar-acquisitions-infographic/. Between 2015 and 2017 alone, these five firms did 175 acquisitions (Gautier and Lamesch, 2020).

[2] See https://www.ftc.gov/news-events/press-releases/2020/02/ftc-examine-past-acquisitions-large-technology-companies. The FTC sued Google and Facebook in October and December 2020.

[3] Note that startup acquisitions are currently virtually unregulated. Indeed, as most of them have a relatively low deal value, they do not need to be reported to antitrust authorities (see Wollmann, 2019).

an incumbent, the incumbent might be able to avoid this outcome by acquiring the startup. However, incumbents are not automatically able to acquire all threatening startups. Instead, their ability to acquire depends on their effort in monitoring the startup scene.

The model reflects the multiple channels through which acquisitions affect innovation and aggregate growth. Some of these channels suggest a negative link. First, incumbents have an incentive to acquire startups in order to preserve their existing profits. However, precisely because they already earn some profits, their marginal benefit from implementing a startup idea is smaller than the one of the startup itself (this is the classical replacement effect first discussed in Arrow, 1962). Thus, some ideas which would have been implemented in the absence of an acquisition might now be shelved. Such events are sometimes called "killer acquisitions" (a term coined by Cunningham, Ederer and Ma, 2020). Second, all else equal, acquisitions slow down creative destruction, by allowing incumbents to avoid displacement more frequently. This creates an economy populated by entrenched incumbents, which have high productivity advantages over their competitors and therefore low innovation incentives.

Other channels instead suggest a positive link between startup acquisitions and growth. First, incumbents might be more efficient at implementing ideas than startups (due to economies of scale and scope, a larger customer base, greater business experience, etc.). When this is the case, acquisitions transfer innovations to more efficient users, and might increase the number of ideas that are successfully implemented.[4] Second, the prospect of an acquisition provides an incentive for startup creation and startup innovation. In our model, acquisitions only occur if the incumbent pays the startup a price that exceeds its outside option of independent entry. Thus, all else equal, incentives for startup creation are higher in the presence of acquisitions. In the business world, many commentators see acquisitions as a natural outcome for startups, and numerous guides advise entrepreneurs how to position their startup in order to be acquired.[5] Finally, startup acquisitions increase the expected lifespan of incumbents, and this increases their innovation incentives. Thus, the effect of lower competition on innovation is actually ambiguous, as famously argued by Aghion, Bloom, Blundell, Griffith and Howitt (2005).[6]

[4]Indeed, this might reflect a beneficial division of labor, with startups specializing in generating ideas and incumbents specializing in implementing them. Likewise, acquisitions might enable startup founders to focus on their core strengths instead of having to deal with management and organizational problems (see https://time.com/3815612/silicon-valley-acquisition for a discussion of these issues).

[5]For some examples, see (1) https://www.forbes.com/sites/alejandrocremades/2019/08/02/how-to-get-your-business-acquired, (2) https://www.inc.com/john-boitnott/how-to-boost-your-businesss-odds-of-an-acquisition or (3) https://thinkgrowth.org/how-to-build-a-startup-that-gets-acquired-85ada592bfd7.

[6]Another effect that we do not explore in our paper is that acquisitions reallocate employees and researchers. The extent of reallocation varies widely: Cunningham *et al.* (2020) show that in the pharmaceutical industry,

To discipline the different forces in our model, we rely on a novel micro dataset for which we combine information on acquisitions (from the ThomsonONE M&A database), patents (from the NBER patent data project) and accounting data (from Compustat). In line with our model, we focus on innovative startups. We define a startup in the data as a firm whose first patent is less than 5 years old. Startups are important for innovation: on average, they represent around 20% of new patents and 60% of new patent citations (indicating that their ideas are on average of higher quality than those of incumbents). Startups are also more likely to be acquired than the average patenting firm.

Our dataset allows us to empirically assess some of the margins through which acquisitions affect innovation. Most importantly, we use it to study the impact of acquisitions on the implementation of ideas. To do so, we analyse the post-acquisition behavior of patent citations. We interpret an increase in the citations received by a startup patent after acquisition as evidence for the associated idea being implemented (consistent with incumbents having an implementation advantage), and a decrease as evidence for the idea being sidelined (consistent with killer acquisitions). To control for selection, we match each patent of an acquired startup to a patent of a non-acquired startup with the same application year, technology class and pre-acquisition citations. For the average industry, we find that acquisition does not affect citations, implying that positive and negative effects roughly cancel out. In some industries, however, killer acquisitions dominate (and in line with Cunningham *et al.*, 2020, this includes the pharmaceutical industry).

While this cross-sectional evidence provides some insights about the effects of acquisitions, it is obviously silent about aggregate general equilibrium effects (such as, for instance, the impact of acquisitions on the startup rate). To take these channels into account, we need to rely on our model. However, we use our cross-sectional findings, as well as other moments from our micro data, to calibrate the model and identify its parameter values.

To understand the link between acquisitions and growth in the calibrated model, we first consider comparative statics with respect to incumbents' startup search costs. These costs can be seen as a reduced-form indicator of frictions in the acquisition market. When they are zero, incumbents may acquire any threatening startup, if they find it optimal to do so. When they are infinite, there are no acquisitions. Our calibration suggests that high search costs (infrequent acquisitions) imply high productivity growth, while low search

only 22% of researchers keep working for the acquirer, but Time Magazine (article cited in Footnote 3) reports this number is three times as high at Google. Tech companies have even coined the term "acqui-hire", with Facebook's CEO Mark Zuckerberg stating that "we have not once bought a company for the company. We buy companies to get excellent people" (https://www.youtube.com/watch?v=OlBDyItD0Ak). Reallocation may change the productivity of the affected researchers (and their colleagues) through knowledge spillovers, discouragement and other effects.

costs (frequent acquisitions) imply low productivity growth.

To understand this relationship, we rely on a useful decomposition. We show that any change in the growth rate from its baseline calibration value can be computed as a weighted average of the change in incumbents' own innovation and the change in innovation due to startup ideas (which, in turn, is the product of the startup rate and the percentage of startup ideas being implemented). The value of the weights is given by our calibration, which targets the share of growth coming from incumbents' own innovation.

Analysing the three sources of variation highlighted by this decomposition, we find that more frequent acquisitions (induced by lower search costs) increase the startup rate, as startups benefit from the option of selling out. However, this slows down creative destruction, so that the average incumbent is more likely to have a high productivity advantage over its competitors, and therefore low incentives to innovate and to implement an acquired idea. Moreover, the higher startup rate erodes the value of incumbents: even though they might buy out many of these startups, they still need to share a part of their rents with them when doing so. As the value of incumbents falls, their innovation incentives decrease further. Therefore, as acquisitions become more frequent, both incumbent innovation and the percentage of implemented startup ideas fall. This more than compensates for the higher startup rate, dragging the growth rate down.

In line with these results, we find that stricter antitrust policy slightly increases growth in our model. In particular, a ban on all startup acquisitions increases the aggregate growth rate by about 0.03 percentage points by year. Again, this occurs despite a significant fall in the startup rate, as the former is compensated by an increase in incumbent innovation and an increase in the percentage of implemented startup ideas.

Finally, we explore how these findings depend on our calibration choices. We find that the negative effects of acquisitions are stronger when we target a higher frequency of startup acquisitions, a higher share of killer acquisitions, and a lower elasticity of substitution between products. This indicates that startup acquisitions are particularly harmful in circumstances where they are frequent, where incumbents do not have large advantages for development costs, and where competition between incumbents is low.

Related literature There is a growing empirical literature on the effect of acquisitions on innovation. The influential recent work of Cunningham *et al.* (2020) on the US pharmaceutical industry provides evidence for several of the channels discussed above. The authors show that acquirers are likely to stop drug research projects of acquired firms when these overlap with their own drug portfolio. These killer acquisitions are more likely if incumbents have a dominant market position. In earlier studies, Seru (2014) and Haucap, Rasch and Stiebale (2019) also provide evidence for a negative effect of mergers and acquisitions (M&As) on firm R&D. Phillips and Zhdanov (2013) instead argue that acquisitions stimulate innovation by small firms that want to be acquired in the future. Using data on publicly traded firms, they show that the R&D of small firms increases after an industry-level acquisition shock. Bena and Li (2014) provide evidence for positive knowledge spillovers after mergers, while Kim (2020) shows that employee mobility after acquisitions can be detrimental to the acquirer in the long run. We provide empirical evidence from a new data set that corroborates some of these findings. However, the main contribution of our paper is to use a general equilibrium model (disciplined by the empirical evidence) to assess the macroeconomic significance of these cross-sectional findings.

On the theoretical side, there has been an intense interest in the industrial organization literature on the effect of M&As on innovation (see Federico, Langus and Valletti, 2017; Cabral, 2018; Bourreau, Jullien and Lefouili, 2018; Bryan and Hovenkamp, 2020; Callander and Matouschek, 2020; Fumagalli, Motta and Tarantino, 2020; Kamepalli, Rajan and Zingales, 2020; Letina, Schmutzler and Seibel, 2020; Denicolò and Polo, 2021). These studies are based on partial equilibrium models, while we take an aggregate general equilibrium perspective.

There are also some recent studies on the macroeconomic effect of M&As. For instance, Dimopoulos and Sacchetto (2017) and David (2020) analyze the effects of M&As on static outcomes such as the allocation of capital, but do not consider innovation and productivity growth.[7] More closely related to us, Cavenaile, Celik and Tian (2020) develop an endogenous growth model with mergers between incumbents, and analyze the effect of these mergers on innovation incentives. Our focus is different, as we study the acquisition of startups by incumbents, leading us to consider novel issues such as killer acquisitions.[8] Finally, Lentz and Mortensen (2016) and Akcigit, Celik and Greenwood (2016) incorporate different versions of a market for ideas (through buyouts or patent sales) in endogenous growth models, showing that such markets improve the allocation of ideas. More broadly, we contribute to the literature on endogenous growth and firm dynamics (Klette and Kortum, 2004; Aghion, Akcigit and Howitt, 2014; Akcigit and Kerr, 2018; Peters, 2020), by extending its standard framework to incorporate acquisitions and study their macroeconomic impact.

[7] There is also an extensive literature on the microeconomic effects of M&As on investment, the allocation of capital, firm productivity and competition. Important studies include Jovanovic and Rousseau (2002), Rhodes-Kropf and Robinson (2008), Blonigen and Pierce (2016), and Wollmann (2019). Some studies have also considered startup acquisitions in particular. For instance, Andersson and Xiao (2016) document a number of stylized facts on startup acquisition in Sweden.

[8] Our paper is also related to Celik, Tian and Wang (2020), who study the effects of information frictions in the merger market on firm innovation and business dynamism.

The remainder of the paper is organized as follows. Section 2 describes our micro-level dataset, and uses it to derive some stylized facts on acquisitions, innovation and the link between these two in the United States. Section 3 presents our model, derives its solution, and describes the main properties of its Balanced Growth Path equilibrium. Section 4 discusses the calibration, comparative statics and our counterfactual experiments. We conclude in Section 5.

2 Data and stylized facts

2.1 Data

Acquisitions data Our data on acquisitions come from the ThomsonONE database. This database provides transaction-level data on mergers and acquisitions (M&As) and includes practically all M&A deals involving US firms between 1981 and 2014. The database provides several variables of interest like the company name, the industry to which the firms belong, the announced and effective dates of the deal, or the transaction value. Revenue and total assets of the parties are also sometimes available. For publicly listed firms, Thomson ONE includes a firm identifier that allows for an easy merge with Compustat and patent data. For private firms, this firm identifier is missing. Consequently, we need to conduct a multi-step cleaning process of private company names before merging the acquisition data to the patent data.

Patent data The patents dataset used in this paper is the NBER Patents Data Project (NBER-PDP), which includes US patent data for 1976-2006.[9]. In addition to the patent owner, this dataset also provides us with the forward and backward citations to the patent, a measure of each patent's originality and generality, and IPC technology classes.

A challenge in matching firm-level data to patents is that firm names are inconsistently recorded on patent files, which leads to many false negative matches. There are two reasons for this: first, the lack of a unique firm identifier in the patent data; second, the lack of uniformity in how company names appear. To address this problem, the NBER-PDP standardizes commonly used words in firm names (Bessen, 2009).

Merging patents to acquisitions The matching procedure between the acquisitions database and the patents database depends on the nature of the firm. For *public* firms,

[9] Source: https://sites.google.com/site/patentdataproject/.

ThompsonONE provides a firm identifier that can be readily matched to NBER-PDP using Compustat as a bridge file in a fairly straightforward exercise. A bigger challenge is the match between patents and the *private* firms. We proceed as follows. ThompsonONE provides us with a company name which we standardize. We then employ a fuzzy name matching algorithm and a large scale manual check to match each company to its patents recorded in the NBER PDP database. As our focus is on acquisitions of startups, which are generally not publicly listed, we have only implemented this matching for private target firms so far.

2.2 Descriptive Statistics

We begin by providing some stylized facts about startup activity and acquisitions in our data. Some of these facts will then be used to calibrate our model.

Importance of startup firms in the economy The three panels in Figure 2 document the importance of startup firms in the acquisition and innovation spheres. We consider a firm to be a startup if it is within 5 years of its first patent application. Panel A displays the share of acquired firms that are defined as startups. Throughout our sample period, out of all acquisitions of private firms with patents in our sample, around 40% are cases where the target firm is a startup (showing that startups are more likely to be bought up than the average patenting firm). The remaining two panels highlight information regarding their innovative ability. From Panel B we learn that startup firms account for 20% of total patent applicants in our sample, and this number is very stable over time. Finally, Panel C reflects the high-quality research done at startups, since 60% of all patent citations are received by the 20% of patents created by startup firms.

Compustat Acquirers Figure 1 compares the size of acquirers to non-acquirers in Compustat. Through our sample period, acquirers are systematically larger than non-acquirers when measured by firm-level sales. This is consistent with the intuition that acquiring firms are not a random sample of Compustat. Rather, they are more likely to be found among the subset of larger Compustat firms.

Comparing acquired to non-acquired startup firms Panel A of Figure 3 shows the startup rate over time. We define the startup rate as the ratio between the number of startups in their first year of existence (i.e., firms that apply for their first patent) to the total number of incumbent patenting firms (i.e., patenting firms whose first patent is more

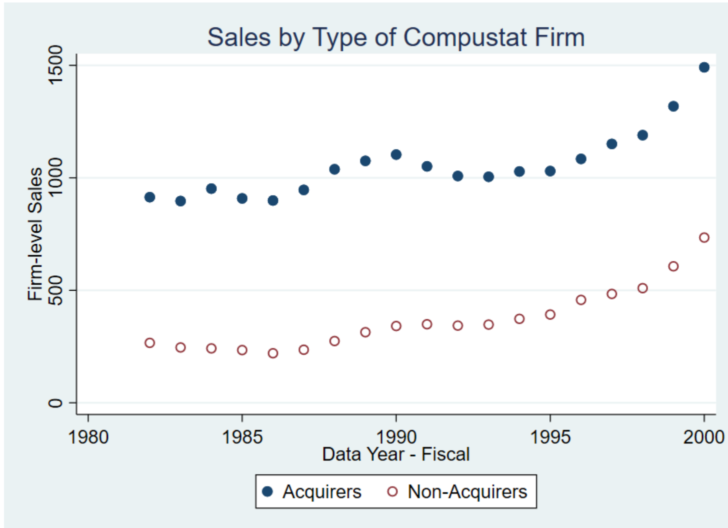

Figure 1: Sales by type of firm: acquirers vs. non-acquirers. Data from ThompsonONE, NBER Patent Data project and Compustat.

than five years old). We can see that this ratio is fairly stable over time at approximately 27%. Panel B illustrates that of the pool of startups present in any given point in time, 1.6% of them will eventually be acquired at some present or future point in time. Finally, Panel C supports the idea that there is also substantial selection on the target firm side. In particular, those 1.6% of acquired startups just mentioned account for a more than proportional share of citations received by their patents (5.5% on average over our sample). This trend is growing over time, however, reaching a peak of 6.5% in the late 1990s.

Mergers and Acquisitions Table C.1 in the Appendix displays the number of M&A transactions taking place between 1980 and 2006, decomposed by whether the acquirer or target firms are public entities or not. Approximately half of the transactions involve a public acquirer, and only around 15% of targets are public. Figure C.1 in the Appendix displays the number of M&A transactions per year and its average transaction value. The observed cyclicality is consistent with previous literature documenting the fact that certain periods experience a boom in acquisitions and sectoral restructurings. This is what we observe in the graph during the late 1980s and the dot-com bubble around the year 2000. The value of these transactions has steadily being rising except for a few years during the 2000s. By 2006, the average transaction value was 400 million US dollars.

Figure C.2 in the Appendix shows the fraction of acquirers and targets with a positive stock of patents at the time of the M&A during the years 1980-2006. For acquirers, this fraction has gone down from around 20% to 13%. The downward trend has been more pronounced for the target companies; while during the mid-1980s this fraction was around

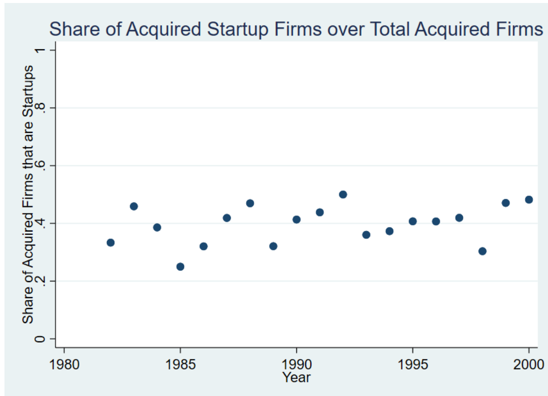

Panel A

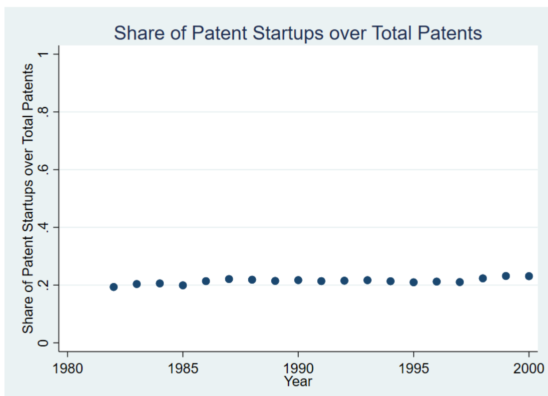

Panel B

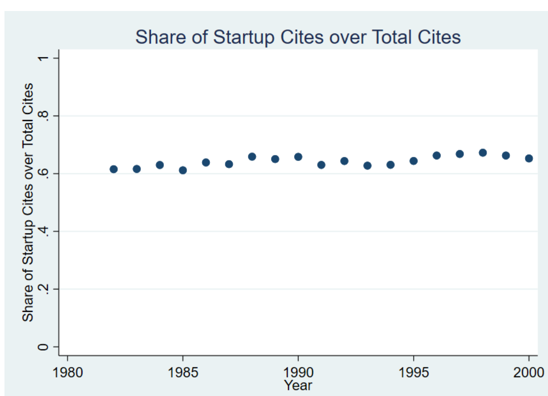

Panel C

Figure 2: Share of acquired startups, share of patents of startups, and citations of startup patents. Data from ThompsonONE, NBER Patent Data project and Compustat (see details in Section 2).

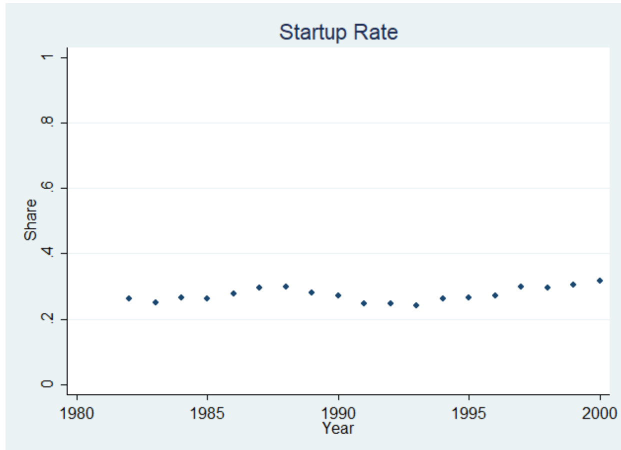

Panel A

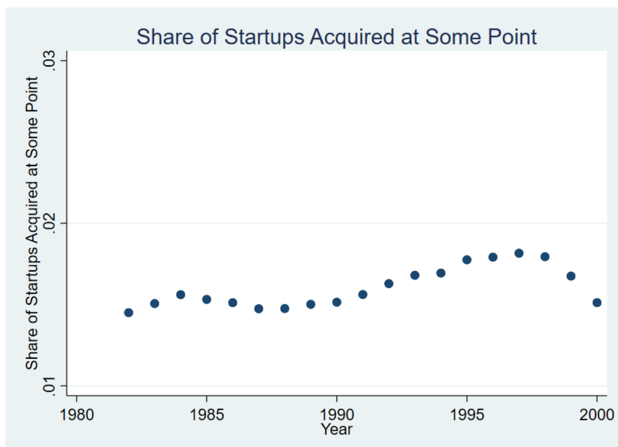

Panel B

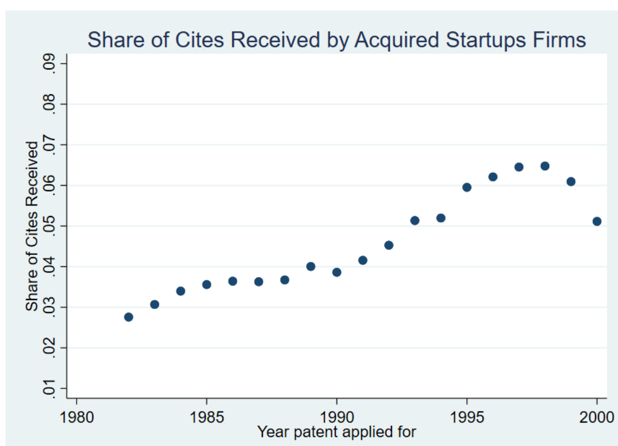

Panel C

Figure 3: Startup rate (by year the patent was applied for), share of startups end up being acquired, and share of startup patents accounted for by startup patent citations. Data from ThompsonONE, NBER Patent Data project and Compustat (see details in Section 2).

20%, it has fallen to 4% by the mid-2000. Finally, Figure C.3 shows the evolution in the number of patents granted to companies involved in M&A transactions, conditional on these firms having a positive stock of patents at the time of the transaction. We see a strong upward trend over the years reaching a point where both acquirer and target have an average stock of patents of 3000 by the year 2006.

2.3 The effect of acquisitions on the implementation of ideas

Our paper aims to assess the net effect of acquisitions on innovation and growth. One important channel linking these phenomena is the effect of acquisitions on the implementation probability of ideas. We therefore use our data to try to assess the causal impact of an acquisition on the probability that a given idea gets implemented: if this is positive, it points towards an implementation advantage of incumbents, if it is negative, it points towards a predominance of killer acquisitions.

Building on previous literature, we proxy whether an idea is implemented by the number of citations it receives after the acquisition, relative to the number of citations received before the acquisition (i.e., we consider the change in patent citations after the acquisition event). Of course, just considering the change in patent citations faces an endogeneity problem: in the previous section, we have shown that acquired patents are different from the average patents. Therefore, we use a matching method (nearest neighbor matching), to link each *treated* patent (i.e., belonging to a startup that will eventually be acquired) to a *control* patent (belonging to a non-acquired startup). We match on several patent and firm characteristics including technological subsector, citations received before acquisition, or patent application year. We artificially assign to each control patent the acquisition year of its matched treated patent.

The regression specification looks as follows:

$$
\begin{aligned}
NumCites_{it} = \beta_0 + \beta_1 \cdot D(Treatment)_i + \beta_2 \cdot D(Post)_{it} \\
+\beta_3 \cdot D(Treatment)_i \cdot D(Post)_{it} + u_{it},
\end{aligned}
$$

where $NumCites_{it}$ are the number of citations received per patent-year, $D(Treatment)_i$ takes value 1 for treated patents, and $D(Post)_{it}$ takes value 1 for the years after acquisition. If $\beta_3 > 0$, then a patent receives more citations (our proxy for the implementation of ideas) after being acquired. Instead, if $\beta_3 < 0$, a patent receives relatively more citations if it is not acquired.

Table 1 presents the estimation results. When a patent changes ownership from a startup to the acquiring firm, its number of citations received (compared to the change

Table 1: Effects of Acquisitions on the Implementation of Ideas

Dep.Var.: Number of Cites	(1)	(2)	(3)	(4)
D(Post)	0.623***	0.496	0.738***	0.503*
	(0.101)	(0.313)	(0.096)	(0.301)
D(Treatment)			0.398***	0.669***
			(0.132)	(0.175)
D(Post)*D(Treatment)	0.145	-0.401**	0.034	-0.337*
	(0.187)	(0.179)	(0.175)	(0.184)
Observations	13,518	1,410	13,536	1,410
Sample	Full	Pharma	Full	Pharma
Year FE	✓	✓	✓	✓
Patent FE	✓	✓		
Matched Pair FE			✓	✓

Notes: A Poisson estimator is used. The dependent variable is the number of citations received at the patent-year level. $D(Treatment)_i$ takes value 1 for treated patents, and $D(Post)_{it}$ takes value 1 for the years after acquisition. significant at 10%; ** significant at 5%; *** significant at 1%.

experienced by the control patent) stays the same for the full sample (columns (1) and (3)), and shrinks by 33% in the pharmaceutical industry (columns (2) and (4)). This latter finding is consistent with the evidence provided by Cunningham *et al.* (2020) for the pharmaceutical industry. Therefore, while killer acquisitions seem to be dominant in this sector, in the average industry negative and positive forces tend to cancel each other out.

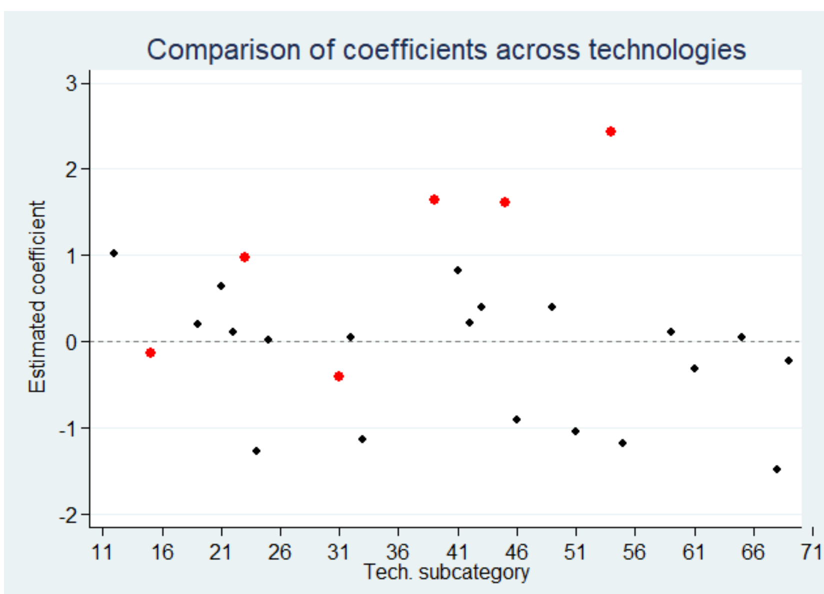

Figure 4: Estimated coefficients of the matching estimation for each technological subcategory

Graphically, Figure 4 displays the estimated coefficient of the interaction term for each technological subcategory. The small blue dots represent non-statistically significant coefficients, while the larger red dots display the ones that are statistically different from zero. The pharmaceutical industry (subcategory 31) is one of the six subcategories with an estimated coefficient that is statistically different from zero. For all the remaining subcategories, we cannot reject a zero value.

This finding indicates that the average acquisition does not seem to have a large effect on the implementation probability of ideas. However, this direct effect is not the only potential link between acquisitions and innovation. To fully study these links in a general equilibrium context, we now introduce our model.

3 Model

In this section, we develop a model of the macroeconomic linkages between startup acquisitions and innovation. While we build on Schumpeterian heterogeneous-firm growth models, our model introduces two important new elements: an distinction between the invention and the implementation of ideas, and the possibility of startup acquisitions.

3.1 Assumptions

Preferences and technology Time is continuous, runs forever and is indexed by $t \in \mathbb{R}_+$. A representative consumer maximizes lifetime utility, given by

$$U = \int_0^{+\infty} e^{-\rho t} \ln\left(C_t\right) \mathrm{d}t, \tag{1}$$

where $\rho > 0$ is the time discount rate and C_t stands for the consumption of the unique final good at instant t. We normalize the price of the final good to one. The household is endowed with L units of time, which she supplies inelastically at the market-clearing wage w_t. Furthermore, the household owns all firms in the economy and accumulates wealth $\mathcal{A}_t$ according to the budget constraint $\dot{\mathcal{A}}_t = r_t \mathcal{A}_t + w_t L - C_t$, where r_t is the rate of return on assets and the price of the final consumption good has been normalized to one.

The final good is produced under perfect competition and assembled from a continuum of differentiated products with a CES production function. Thus, final output is

$$Y_t = \left(\int_0^1 \left(\omega_{jt}\right)^{\frac{1}{\varepsilon}} \left(y_{jt}\right)^{\frac{\varepsilon-1}{\varepsilon}} \mathrm{d}j \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{2}$$

where y_{jt} is the output of product j at instant t, ω_{jt} is the quality of product j at instant t, and $\varepsilon > 1$ is the elasticity of substitution between products. Product quality follows an exogenous stochastic process. We assume that quality can take values in a finite set Ω, and that firms transition from state ω to state ω' at a Poisson rate $\tau_{\omega,\omega'}$. We also assume that the economy starts in the steady state of this process, and for convenience we normalize $\int_0^1 \omega_{jt} \mathrm{d}j = 1$.

Each product can potentially be produced by a large number of firms f, with a linear production technology using labor:

$$y_{jft} = a_{jft} l_{jft}, \tag{3}$$

where y_{jft} is the output of product j by firm f at instant t, a_{jft} is the productivity of the firm, and l_{jft} is the labor input. We assume that there is static Bertrand competition on product markets. As we show later, this implies that each product is only produced by the highest-productivity firm in equilibrium. We denote the productivity of this firm by a_{jt}, and define average productivity A_t as

$$A_t \equiv \left(\int_0^1 a_{jt}^{\varepsilon-1} dj \right)^{\frac{1}{\varepsilon-1}}. \tag{4}$$

Productivity is improved through innovations, which are the result of a two-step process. First, firms invest into research in order to generate new ideas. Then, they invest into development in order to implement these ideas and turn them into innovations. The next sections describe these research and development (R&D) technologies.

Research and Development Innovations are generated by incumbent firms (i.e., firms which already produce at instant t) and by a large mass of potential entrants, which we refer to as startups.

To generate an idea at a Poisson arrival rate z, an incumbent must pay a research cost of $\xi_I \cdot z^{\psi} \cdot \widetilde{a}_{jt}^{\varepsilon-1} \cdot Y_t$ units of the final good. In this cost function, $\xi_I > 0$ is a cost shifter, $\widetilde{a}_{jt} \equiv \frac{a_{jt}}{A_t}$ is the relative productivity of the incumbent firm, and $\psi > 1$ is the elasticity of research output with respect to research spending. Thus, research costs are increasing and convex in the arrival rate of ideas. Furthermore, they are proportional to the relative productivity of the incumbent and to aggregate GDP. These scaling assumptions are necessary to ensure balanced growth.[10]

[10] In particular, the fact that costs scale with relative productivity makes research choices independent of current productivity, as in Peters (2020). Without this assumption, more productive incumbents innovate more, and production is eventually taken over by an arbitarily small number of firms.

To implement an idea, the incumbent needs to invest into development. Precisely, if the incumbent invests $\kappa_I \cdot i_I^{\psi} \cdot \widetilde{a}_{jt}^{\varepsilon-1} \cdot Y_t$ units of the final good (with $\kappa_I > 0$), it successfully implements the idea with probability i_I.[11] As usual in endogenous growth models, we assume that productivity evolves on a ladder, with step size $\lambda > 1$. An implemented idea (an innovation) increases the productivity of the incumbent by one step on this ladder, i.e., by a factor λ. Instead, an idea that is not implemented disappears forever. Therefore, ideas are either implemented immediately or never.

Ideas and innovations are also generated by startups. We assume that a startup can be created at a fixed cost $\xi_S \cdot Y_t$, and generates a Poisson arrival rate 1 of ideas. A startup's idea applies to a randomly drawn good $j \in [0, 1]$. As for incumbents, startup ideas are either implemented immediately or never. Precisely, when the startup invests $\kappa_S \cdot i_S^{\psi} \cdot \widetilde{a}_{jt}^{\varepsilon-1} \cdot Y_t$ units of the final good (with $\kappa_S > 0$), it implements the idea with probability i_S. To reflect the empirical fact that startup ideas might represent larger advances than incumbent ideas, we assume that a startup idea increases productivity by $n_S = 1 + N$ steps (of size λ each), where $N \in \mathbb{N}$ is drawn from a Poisson distribution with parameter γ. Thus, on average, a startup idea represents γ more steps on the productivity ladder than an incumbent idea. Importantly, we assume that the quality of the idea is only revealed after investing into development.

In equilibrium, a startup that implements its idea displaces the incumbent producer of product j and becomes the new incumbent in this product line. However, the startup may not always choose to implement: alternatively, it can be acquired by the incumbent. In the next section, we describe these acquisitions.

Acquisitions We assume that acquisitions can take place if, and only if, there is a "meeting" between the startup and the threatened incumbent producer.

The meeting probability is endogenous, and depends on the effort of the incumbent in monitoring the startup scene. We assume an incumbent needs to spend $\chi \cdot s^{\varphi} \cdot \widetilde{a}_{jt}^{\varepsilon-1} \cdot Y_t$ (with $\chi > 0$ and $\varphi > 1$) units of the final good in order to generate a probability s to meet a startup that innovates on its product. Thus, the search costs for startups are increasing and convex in the search effort. As usual, they also scale with relative productivity and aggregate GDP to ensure balanced growth. We think of this framework as a reduced-form model of information and search frictions in the acquisition market. These frictions prevent incumbents from noticing all threatening startups and force them to spend resources in

[11] In fact, we assume that the implementation probability is given by $\min(i_I, 1)$, so that it is always well defined. However, we choose parameter values ensuring that firms never choose an implementation probability of 1. For simplicity, we therefore omit the min operator in the text. The same statement applies to all other implementation and meeting probabilities introduced below.

order to monitor the market.

When there is a meeting, the incumbent may acquire the startup. The incumbent then transfers p^A_{jt} units of the final good (the acquisition price) to the startup, in exchange for the startup exiting forever and handing over its idea to the incumbent. The incumbent then invests into the development of the startup's idea, using its own development technology. That is, by investing $\kappa_I \cdot i_A^{\psi} \cdot \widetilde{a}_{jt}^{\varepsilon-1} \cdot Y_t$ units of the final good, it implements the idea with probability i_A.

Acquisitions occur if, and only if, they generate a surplus, that is, if and only if the joint value of both firms after the acquisition is larger than the sum of their outside options. The acquisition price is determined through Nash bargaining over the surplus, where the incumbent has a bargaining weight $\alpha \in (0, 1)$. There are two reasons for which acquisitions may generate a surplus in the model. First, the startup's idea may be more valuable in the hands of the incumbent (e.g., because the latter has lower development costs). Second, acquisitions prevent entry, and therefore prevent the destruction of incumbent rents. While the first force corresponds to a socially valuable transfer of ideas, the second does not. As we will see later, the relative strength of these forces plays an important role for the aggregate implications of startup acquisitions.

Timing Figure 5 summarizes the timing of events for a startup idea within an instant of length $(t, t + dt)$. After the idea appears, the incumbent might or might not notice it, depending on the search probability s. If there is no meeting, the startup decides whether or not to implement it (with probability i_S), which leads to possible entry and displacement of the incumbent. If there is a meeting, then there is an acquisition if, and only if, the acquisition surplus is positive. In case of an acquisition, the incumbent then chooses the probability i_A with which to implement the startup's idea.

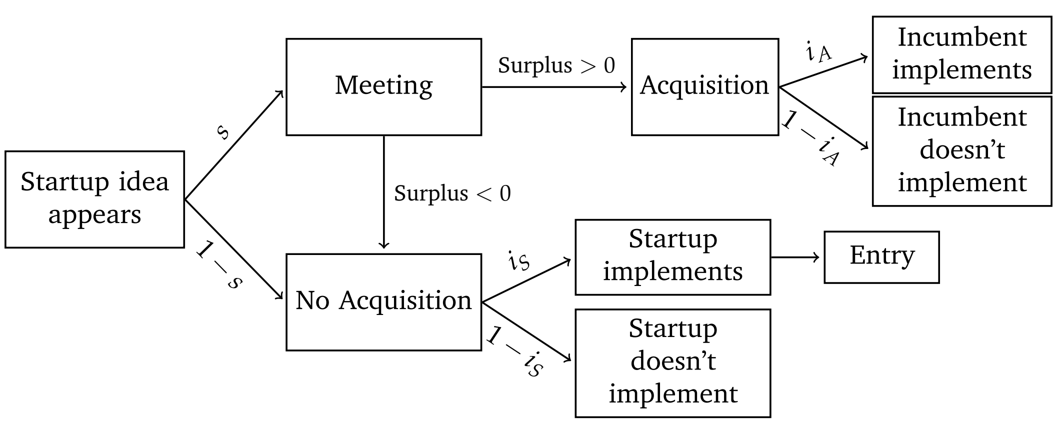

Figure 5: Timing of events for a startup idea within a period $(t, t + dt)$.

3.2 Equilibrium

Throughout, we consider a balanced growth path (BGP) equilibrium with positive entry, in which all aggregate variables grow at a constant rate g.

3.2.1 Household decisions, prices and profits

On the BGP, the representative consumer's optimal consumption choice satisfies the Euler equation

$$\frac{\dot{C}_t}{C_t} \equiv g = r - \rho. \tag{5}$$

Bertrand competition implies that each product is only produced by the highest-productivity firm. However, pricing decisions depend on the relative productivity of this firm with respect to its closest follower (the firm with the second-highest productivity). As productivity evolves on a ladder, we can define the "technology gap" (the number of productivity steps between the incumbent and the follower), as the integer n_{jt} holding

$$\lambda^{n_{jt}} \equiv \frac{a_{jt}}{a_{jt}^F}, \tag{6}$$

where a_{jt}^F is the productivity of the follower. Note that in our model, the follower is an old incumbent: once a startup displaces an incumbent, the latter becomes the new follower.

The demand for each product j is given by the isoelastic function $y_{jt} = \omega_{jt} \cdot \left(p_{jt}\right)^{-\varepsilon} \cdot Y_t$. Thus, if incumbents could freely choose their price, they would set a constant markup over their marginal cost. However, their price must also be low enough to keep the follower out of the market. For any product j, the average cost of the follower at instant t is by a factor $\lambda^{n_{jt}}$ higher than the one of the incumbent. Thus, when the incumbent charges a markup $\lambda^{n_{jt}}$, the follower makes zero profits and does not produce. Accordingly, markups are

$$\mu(n_{jt}) = \min\left(\lambda^{n_{jt}}, \frac{\varepsilon}{\varepsilon - 1}\right). \tag{7}$$

For high technology gaps, the incumbent can charge the monopoly markup, while for low technology gaps, it must charge a lower markup to keep the follower out.

This markup choice implies that the price of any product j is given by $p_{jt} = \mu(n_{jt}) \cdot \frac{w_t}{a_{jt}}$, and profits are

$$\pi_t\left(\omega_{jt}, n_{jt}, a_{jt}\right) = \omega_{jt} \cdot \left(1 - \frac{1}{\mu\left(n_{jt}\right)}\right) \cdot \left(\mu\left(n_{jt}\right)\right)^{1-\varepsilon} \cdot \left(\frac{a_{jt}}{w_t}\right)^{\varepsilon - 1} \cdot Y_t \tag{8}$$

Equation (8) shows that profits are increasing in product quality ω_{jt}, in productivity a_{jt} and in the technology gap n_{jt}. In particular, note that profits are concave in n_{jt}. Indeed, higher technology gaps imply higher markups, but as the firm approaches the unconstrained monopoly markup, these gains become smaller and eventually vanish.

3.2.2 Research, Development and Acquisitions

Incumbent's dynamic decisions At every point in time, incumbents need to choose an optimal level of research spending z and search effort s. Moreover, whenever they obtain an idea, they need to choose an optimal level of development spending, and whenever they meet a startup, they must decide whether to acquire it.

The dynamic problem of the incumbent has two endogenous state variables (the technology gap n and productivity a) and one exogenous state variable (product quality ω). Furthermore, the value function also depends on some aggregate variables, which change over time. Thus, we denote the value function by $V_t(\omega, n, a)$. On the BGP, the Hamilton-Jacobi-Bellman (HJB) equation is

$$\begin{aligned} r \cdot V_t(\omega,n,a) = \max_{z,s} \Bigg\{ & \underbrace{\pi_t(\omega,n,a)}_{\text{Profits}} - \underbrace{\xi_I \cdot z^{\psi} \cdot \widetilde{a}_t^{\varepsilon-1} \cdot Y_t}_{\text{Research cost}} - \underbrace{\chi \cdot s^{\varphi} \cdot \widetilde{a}_t^{\varepsilon-1} \cdot Y_t}_{\text{Search effort}} \\ & + \underbrace{z \cdot \max_{i_I} \left[i_I \cdot \Big(V_t(\omega, n+1, \lambda a) - V_t(\omega,n,a) \Big) - \kappa_I \cdot i_I^{\psi} \cdot \widetilde{a}_t^{\varepsilon-1} \cdot Y_t \right]}_{\text{Own innovation}} \\ & + \underbrace{x \cdot \left[s \cdot V_t^{\text{Meet}}(\omega,n,a) + (1-s) \cdot V_t^{\text{NoMeet}}(\omega,n,a) - V_t(\omega,n,a) \right]}_{\text{Startup appears}} \Bigg\} \\ & + \underbrace{\sum_{\omega' \in \Omega} \tau_{\omega,\omega'} \cdot \Big[V_t(\omega',n,a) - V_t(\omega,n,a) \Big]}_{\text{Quality shock}} + \underbrace{\dot{V}_t(\omega,n,a)}_{\text{Drift}}. \end{aligned} \tag{9}$$

The HJB equation shows how the discounted value of the firm changes over time. First, at every instant, the firm collects static profits and spends on research and startup search, as shown in the first line. As shown in the second line, the incumbent discovers an idea at Poisson rate z, and then chooses the optimal development investment i_I. An implemented idea increases its technology gap by one step and its productivity by a factor of λ. The third line shows that at rate x, a startup makes an innovation on the incumbent's product. In that case, there is a meeting (and thus potentially an acquisition) with probability s, and no meeting with probability $1 - s$. We denote by $V_t^{\text{Meet}}(\omega, n, a)$ the expected continuation value of the incumbent in case there is a meeting, and by $V_t^{\text{NoMeet}}(\omega, n, a)$ the expected continuation value of the incumbent in case there is no meeting. Finally, the fourth line shows that the incumbent is subject to exogenous product quality shocks, and that its value drifts over time due to aggregate growth.

Acquisitions and Startup creation To analyze the interaction between an incumbent and a startup that threatens to replace it, we first consider the case in which there is no meeting between both firms. In that case, there is no acquisition, and the incumbent's expected continuation value is

$$V_t^{\text{NoMeet}}(\omega, n, a) = \Big[1 - i_{S,t}(\omega, n, a)\Big] \cdot V_t(\omega, n, a), \tag{10}$$

where $i_{S,t}(\omega, n, a)$ is the optimal development investment of a startup facing an incumbent of type (ω, n, a). When the startup does not implement its idea, the incumbent's continuation value is just its current value. Instead, when the startup implements its idea, the incumbent is displaced and its continuation value is zero.

Likewise, we can derive the expected value of a startup in the absence of a meeting, denoted by $V_{S,t}^{\text{NoMeet}}(\omega, n, a)$. This quantity holds

$$V_{S,t}^{\text{NoMeet}}(\omega, n, a) = \max_{i_S} \left\{ i_S \cdot \left(\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot V_t(\omega, n_S, \lambda^{n_S} a) \right) - \kappa_S \cdot i_S^{\psi} \cdot \widetilde{a}_t^{\varepsilon-1} \cdot Y_t \right\} \tag{11}$$

where $\theta(n_S) \equiv e^{-\gamma} \cdot \frac{\gamma^{n_S-1}}{(n_S-1)!}$ denotes the probability that the startup's innovation advances productivity by $n_S = 1, 2, \ldots$ steps. In the absence of a meeting, a startup chooses an optimal level of development investment i_S. When its idea is implemented, the startup becomes the new incumbent producer. With probability $\theta(n_S)$, it takes n_S steps on the productivity ladder. It then has a technology gap of n_S (over the previous incumbent, which is now the follower) and productivity $\lambda^{n_S} a$. On the other hand, if the idea fails, the startup exits forever and has a continuation value of zero.

Next, we turn to the case in which a meeting does take place. To determine whether this leads to an acquisition, we compute the surplus that would be generated by an acquisition, denoted $\Sigma_t(\omega, n, a)$. The surplus holds

$$\begin{aligned}\Sigma_t(\omega,n,a) = \max_{i_A} \Bigg\{ &(1-i_A)\cdot V_t(\omega,n,a) + i_A \cdot \sum_{n_S=1}^{+\infty} \theta(n_S)\cdot V_t(\omega, n+n_S, \lambda^{n_S} a) \\ &- \kappa_I \cdot i_A^{\psi} \cdot \widetilde{a}_t^{\varepsilon-1} \cdot Y_t \Bigg\} - V_t^{\text{NoMeet}}(\omega,n,a) - V_{S,t}^{\text{NoMeet}}(\omega,n,a)\,. \end{aligned} \tag{12}$$

In equation (12), the term inside the curly brackets captures the joint value of incumbent and startup after an acquisition. The acquisition allows the incumbent to keep its baseline value $V_t(\omega, n, a)$. Moreover, the incumbent acquires the startup's idea and chooses an optimal development investment i_A in order to implement it. In case of success, the quality of the idea is revealed, and an idea of quality n_S improves the incumbent's technology gap by n_S units and its productivity by a factor λ^{n_S}. Finally, the incumbent transfers the acquisition price to the startup (and this acquisition price is the startup's post-acquisition value). As this is a pure transfer, it does not feature in the joint value shown above. To obtain the surplus, we subtract from the joint value the outside options of incumbent and startup, which are equal to their expected values in the absence of a meeting.

An acquisition takes place if, and only if, the expected surplus is positive. Then, the surplus is split between both firms according to their Nash bargaining weights. Accordingly, the continuation value for an incumbent in case of a meeting with the startup is:

$$V_t^{\text{Meet}}(\omega,n,a) = V_t^{\text{NoMeet}}(\omega,n,a) + \alpha \cdot \max\Big(0, \Sigma_t(\omega,n,a)\Big). \tag{13}$$

For the startup, the continuation value conditional on meeting the incumbent is

$$V_{S,t}^{\text{Meet}}(\omega,n,a) = V_{S,t}^{\text{NoMeet}}(\omega,n,a) + (1-\alpha)\cdot \max\Big(0, \Sigma_t(\omega,n,a)\Big). \tag{14}$$

Whenever an acquisition takes place, this continuation value is also equal to the acquisition price. Finally, in an equilibrium with positive startup creation ($x > 0$), the following free-entry condition must hold:

$$\xi_S \cdot Y_t = \mathbb{E}_t\Big[s_t(\omega,n,a)\cdot V_{S,t}^{\text{Meet}}(\omega,n,a) + \Big(1 - s_t(\omega,n,a)\Big)\cdot V_{S,t}^{\text{NoMeet}}(\omega,n,a)\Big]. \tag{15}$$

where $s_t(\omega, n, a)$ denotes the optimal search effort by the incumbent. This equation shows that the cost of creating a startup, $\xi_S \cdot Y_t$, must be equal to the expected benefit of creating a startup, shown on the right-hand side. The startup's idea falls on a randomly chosen product j, characterized by a quality ω, a technology gap n and productivity a. The expectation operator refers to the joint distribution of products over these states. Depending on whether the startup meets an incumbent or not, it then obtains one the continuation values defined in equations (11) and (14).

Optimal policies To solve for the BGP policies, we guess and verify that the incumbent's value function holds $V_t(\omega, n, a) = v(\omega, n) \cdot \widetilde{a}_t^{\varepsilon-1} \cdot Y_t$, i.e., that the value function scales (with a time-invariant factor of proportionality) in relative productivity and aggregate GDP. Furthermore, we conjecture that aggregate productivity A_t grows at the same rate as aggregate consumption C_t and wages w_t. These guesses allow us to simplify the dynamic problem considerably. First, combining them with the Euler equation (5) and the continuation values defined in equations (10) and (13), we can rewrite the HJB equation as

$$\begin{aligned}\Big(\rho + (\varepsilon - 1)\, g\Big) \cdot v(\omega, n) = \max_{z,s} \Bigg\{ & \omega \cdot \left(1 - \frac{1}{\mu\,(n)}\right) \cdot (\mu\,(n))^{1-\varepsilon} \cdot \left(\frac{A_t}{w_t}\right)^{\varepsilon-1} - \xi_I \cdot z^{\psi} - \chi \cdot s^{\varphi} \\ & + z \cdot \max_{i_I} \left[i_I \cdot \Big(\lambda^{\varepsilon-1} \cdot v(\omega, n+1) - v(\omega, n)\Big) - \kappa_I \cdot i_I^{\psi} \right] \\ & + x \cdot \Big[s \cdot \alpha \cdot \widetilde{\sigma}(\omega, n) - i_S(\omega) \cdot v(\omega, n) \Big] \Bigg\} \\ & + \sum_{\omega'} \tau_{\omega,\omega'} \cdot \Big[v(\omega', n) - v(\omega, n) \Big] \end{aligned} \tag{16}$$

where $\widetilde{\sigma}(\omega, n) \equiv \frac{\max(0, \Sigma_t(\omega,n,a))}{\widetilde{a}_t^{\varepsilon-1} \cdot Y_t}$, the normalized acquisition surplus, is time-invariant as shown in Appendix A.1. This equation pins down the value function v as a function of three endogenous aggregate constants: aggregate growth g, the startup rate x and the productivity-to-wage ratio $\frac{A_t}{w_t}$.[12]

The HJB equation implies that the incumbent's optimal research investment is

$$z(\omega, n) = \left[\frac{i_I(\omega, n) \cdot \Big(\lambda^{\varepsilon-1} \cdot v(\omega, n+1) - v(\omega, n)\Big) - \kappa_I \cdot \big(i_I(\omega, n)\big)^{\psi}}{\xi_I \psi} \right]^{\frac{1}{\psi - 1}} \tag{17}$$

where $i_I\,(\omega, n)$ is the optimal development investment chosen by the incumbent for its own ideas. As usual, the firm equalizes the marginal cost of research to its marginal benefit, which is the arrival of an undeveloped idea.

[12] As each startup has a Poisson arrival rate 1 of ideas, x corresponds both to the mass of startups and the arrival rate of startup ideas. As there is a mass 1 of incumbents, x is also the startup rate. Note that because average productivity and wages grow at the same rate, A_t/w_t is a constant.

In turn, the optimal startup search investment is given by

$$s(\omega,n) = \left(\frac{x \cdot \alpha \cdot \widetilde{\sigma}(\omega,n)}{\chi\varphi}\right)^{\frac{1}{\varphi-1}}. \tag{18}$$

Intuitively, the search effort is increasing in the arrival rate of startup ideas x, in the acquisition surplus $\widetilde{\sigma}(\omega,n)$ and in the incumbent's surplus share α.

Regarding development, the investment of incumbents into their own ideas holds

$$i_I(\omega,n) = \left(\frac{\lambda^{\varepsilon-1} \cdot v(\omega,n+1) - v(\omega,n)}{\kappa_I\psi}\right)^{\frac{1}{\psi-1}}. \tag{19}$$

Again, firms equalize the marginal cost of development to its marginal benefit, which comes from improving productivity and widening the technology gap.

Investment of incumbents into acquired ideas holds

$$i_A(\omega,n) = \left(\frac{\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot \lambda^{n_S\cdot(\varepsilon-1)} \cdot v(\omega,n+n_S) - v(\omega,n)}{\kappa_I\psi}\right)^{\frac{1}{\psi-1}} \tag{20}$$

Finally, the optimal development investment by startups, defined in equation (11), is given by

$$i_S(\omega) = \left(\frac{\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot \lambda^{n_S\cdot(\varepsilon-1)} \cdot v(\omega,n_S)}{\kappa_S\psi}\right)^{\frac{1}{\psi-1}} \tag{21}$$

Comparing equation (20) with equation (21) shows that incumbents and startups may make different development choices for the same idea. These differences stem from three sources. First, development costs may be different, and all else equal, a lower marginal cost (a lower cost shifter κ) implies higher investment. Second, incumbents can apply their innovation to their existing high technology gap (and as the profit function (8) shows, productivity and markup are complements). Third, as the value function is concave in the technology gap n, there is an Arrow replacement effect: the fact that the incumbent already earns some monopoly rents makes it less attractive to implement. When this last effect dominates, some ideas that would have been implemented by a startup will be shelved by the incumbent (i.e., some acquisitions will be killer acquisitions).

Finally, as shown in greater detail in Appendix A.1, our guesses imply that the quality, productivity and technology gap of a product at at a given point in time are independent random variables. Therefore, we have

$$m_t(\omega, n, a) = m(\omega) \cdot m(n) \cdot m_t(a), \tag{22}$$

where $m(\bullet)$ stands for the mass of products with a certain characteristic. Using this property, the free entry condition simplifies to

$$\xi_S = \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \cdot \left[v_S^{\text{NoMeet}}(\omega, n) + s(\omega, n) \cdot (1 - \alpha) \cdot \widetilde{\sigma}(\omega, n) \right]. \tag{23}$$

This shows that research, development and acquisition decisions are independent of productivity a. Therefore, we do not need to keep track of the productivity distribution. The invariant distribution of quality ω is exogenous. Finally, the invariant distribution of technology gaps n depends on innovation and acquisition decisions, and is derived in Appendix A.2.

In equilibrium, the startup rate x will be such that the expected value of startup creation equals the fixed cost ξ_S. However, as the previous equations show, firms' innovation and acquisition decisions also depend on two other aggregate variables, the productivity-to-wage ratio $\frac{A_t}{w_t}$ and aggregate growth g. To close the model, we derive these variables in the next section.

3.2.3 Closing the model

First, using the definition of the CES price index, in Appendix A.3 we show that the productivity-to-wage ratio holds

$$\frac{A_t}{w_t} = \left(\sum_{n=1}^{+\infty} m(n) \cdot \big(\mu(n)\big)^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}}. \tag{24}$$

This shows that, along the BGP, aggregate productivity A_t grows at the same rate as the wage w_t. The ratio of both variables depends on the markup distribution across incumbents.

Next, we need to impose labor market clearing. Labor is in fixed supply $L > 0$ and is used only in production. Imposing labor market clearing gives an expression for the aggregate labor share (see details in Appendix A.3):

$$\frac{w_t L}{Y_t} = \left(\frac{A_t}{w_t} \right)^{\varepsilon - 1} \sum_{n=1}^{+\infty} m(n) \cdot \big(\mu(n)\big)^{-\varepsilon}. \tag{25}$$

Product market clearing, in turn, implies that aggregate output is fully used for consumption (C_t), research (R_t), development (D_t) and search (S_t). Therefore, we have

$$Y_t = C_t + R_t + D_t + S_t \tag{26}$$

where

$$R_t = Y_t \cdot \left(\sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \cdot \xi_I \cdot \big(z(\omega, n)\big)^{\psi} + x \cdot \xi_S \right) \tag{27}$$

$$\begin{aligned} D_t = Y_t \cdot \Bigg(\sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \cdot \Big[& z(\omega, n) \cdot \kappa_I \cdot \big(i_I(\omega, n)\big)^{\psi} \\ & + x \cdot \Big(\widetilde{s}(\omega, n) \cdot \kappa_I \cdot \big(i_A(\omega, n)\big)^{\psi} + \big(1 - \widetilde{s}(\omega, n)\big) \cdot \kappa_S \cdot \big(i_S(\omega)\big)^{\psi} \Big) \Big] \Bigg), \end{aligned} \tag{28}$$

$$S_t = Y_t \cdot \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \cdot \chi \cdot \big(s(\omega, n)\big)^{\varphi}, \tag{29}$$

where $\widetilde{s}(\omega, n) \equiv s(\omega, n) \cdot \mathbb{1}_{\widetilde{\sigma}(\omega,n)>0}$ is the probability that, conditional on an arrival of a startup idea on a product of type (ω, n), an acquisition occurs. This shows that consumption grows at the same rate as output.

Finally, as shown in Appendix A.4, the growth rate is

$$g = \frac{1}{\varepsilon - 1} \cdot \left[\sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \cdot b(\omega, n) \right], \tag{30}$$

where $b(\omega, n) \equiv b_I(\omega, n) \cdot \left(\lambda^{\varepsilon-1} - 1\right) + b_S(\omega, n) \cdot \left(\lambda^{\varepsilon-1} \cdot e^{\gamma \cdot \left(\lambda^{\varepsilon-1} - 1\right)} - 1\right)$ is the overall arrival rate of innovations (on average across innovation steps), and

$$\begin{aligned} b_I(\omega, n) &\equiv z(\omega, n) \cdot i_I(\omega, n) \\ b_S(\omega, n) &\equiv x \cdot \Big(\widetilde{s}(\omega, n) \cdot i_A(\omega, n) + \big(1 - \widetilde{s}(\omega, n)\big) \cdot i_S(\omega) \Big) \end{aligned}$$

are the arrival rates of innovations generated by incumbents (b_I) and startups (b_S), respectively. The formula illustrates that when γ is positive, (implemented) startup ideas contribute relatively more to growth.

This concludes the description of our model's equilibrium conditions. Appendix B provides details on its numerical solution. In the next section, we proceed to analyse its quantitative implications.

4 Quantitative Analysis

Our quantitative analysis proceeds in several steps. First, in Section 4.1, we calibrate our model's parameters, matching aggregate and micro-level moments (including several moments introduced in Section 2). In Section 4.2, we review some relevant qualitative properties of the model, and in Section 4.3 we study the relationship between the frequency of acquisitions and aggregate growth. In Section 4.4, we evaluate the effects of antitrust policy, and in Section 4.5, we discuss robustness checks.

4.1 Calibration Strategy

We assume that a period of length 1 in the model corresponds to one year in the data. Then, we set several parameters externally. First, we set the discount rate to $\rho = 0.02$ (which, combined with a 2% growth rate, implies a 4% annual interest rate). Second, we set the elasticity of substitution to $\varepsilon = 4$, a standard value in the literature (Aghion, Bergeaud, Boppart, Klenow and Li, 2021; Galí and Monacelli, 2016). Section 4.5 contains a robustness exercise showing that this parameter is important for the magnitude of our results.

We assume that there are two product quality classes, $\Omega = \{\omega_L, \omega_H\}$ with $\omega_L < \omega_H$. At every point in time, 50% of firms belong to the H class, and their sales account for 94% of GDP (in line with the average industry-level sales share of firms with above-median sales in Compustat). This implies that $\omega_H/\omega_L = 15.7$. Moreover, we assume that firms transition from ω_H to ω_L at a Poisson rate $\tau = 0.11$, matching the fact that in every year, 11% of Compustat firms above the sales median of their industry drop below this sales median. We set the elasticity of R&D costs to innovation to $\psi = 2$, following empirical evidence in Akcigit and Kerr (2018). Following David (2020), we set the Nash bargaining parameter for incumbents to $\alpha = 0.5$. Finally, we set the average step size advantage for startup ideas to $\gamma = 0.36$. To obtain this number, we rely on findings from Kogan, Papanikolaou, Seru and Stoffman (2017), which estimates that the elasticity of a patent's market value to the number of its forward citation 0.17. Our results from Section 2 indicate that the average startup patent is cited six times as much as the average incumbent patent. Therefore, we assume that a startup patent represents on average $\gamma = 6^{0.17} - 1 \approx 36\%$ more steps than an incumbent patent.

This leaves seven parameters to be identified: the productivity step size, λ; the research and development cost shifters for incumbents, ξ_I and κ_I; the fixed cost of startup creation, ξ_S; the development cost shifter for startups, κ_S; and the scale and curvature parameters in the incumbent's effort cost function, χ and φ. We calibrate these parameters internally using an indirect inference approach: that is, we choose the set of parameter values that minimizes the distance between a set of model-generated moments and their empirical counterparts.[13] The success of this calibration strategy relies on choosing moments that are both relevant for the economic intuitions we want to highlight, as well as sufficiently sensitive to variation in individual parameters. As the model is highly non-linear, however, all moments are affected by all parameters to various degrees, making identification challenging. Nevertheless, we can provide economic intuitions for how some moments have more identification power than others. To support these intuitions, in Appendix B.2 we perform a more rigorous global identification exercise, and demonstrate that all parameters are well-identified by our chosen set of moments.

Table 2 lists the calibrated parameter values and summarizes the model fit. First, we target a growth rate of 2%, consistent with the long-run growth rate of GDP per capita in the United States (Jones, 2016). This moment identifies the innovation step λ, and we find that an innovation roughly increases productivity by 4%.

Second, we target the startup rate (defined in the data as the ratio between the number of startups and the total number of patenting firms, see Panel A of Figure 3). In our model, this corresponds to the mass of startups x (as the mass of incumbents is 1 by definition). This moment identifies ξ_S, the cost of startup creation.

Third, we target the percentage of startups that are acquired. In the data, this percentage equals 5.5% (when weighting by patent citations, as we do to account for the fact that startups are heterogeneous in the data, but homogeneous in our model).[14] This moment identifies χ, the search cost of incumbents for startups.

Forth, we set the development cost scale parameters of incumbents (κ_I) and startups (κ_S) to target their average implementation probabilities of startup ideas. These implementation probabilities are not directly observed in the data, but we can infer them indirectly. Precisely, in our model, a startup faces three possible outcomes: it can be acquired, it can innovate and enter, or it can fail in its innovation effort. As explained above, we target 5.5% of all startups to be acquired. Moreover, as we observe an entry rate of 9.4% in the data, the

[13]Formally, the vector of parameters $\boldsymbol{\theta} = (\lambda, \xi_I, \kappa_I, \xi_S, \kappa_S, \chi, \varphi)$ is chosen to minimize the following criterion distance function: $\sum_{m=1}^{M} \frac{|\text{Moment}_m(\text{Model},\boldsymbol{\theta}) - \text{Moment}_m(\text{Data})|}{0.5|\text{Moment}_m(\text{Model},\boldsymbol{\theta})| + 0.5|\text{Moment}_m(\text{Data})|}$.

[14]In the model, the percentage of acquired startups equals

$$ShareStartupsAcq \equiv \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \cdot \widetilde{s}(\omega, n).$$

Table 2: Calibrated parameters and model fit.

A. Externally Calibrated Parameters

Parameter	*Description*	*Value*	*Target/Source*
ρ	Discount rate	0.02	4% annual real interest rate
ε	Elasticity of substitution	4	Standard in the literature
ω_H/ω_L	Relative product quality	15.7	Top 50% sales share (Compustat)
τ_{HL}	Transition rate from high to low quality	0.11	Likelihood to drop from Top 50% (Compustat)
ψ	R&D cost curvature	2	Akcigit and Kerr (2018)
α	Bargaining weight for incumbents	0.5	David (2020)
γ	Step size advantage of startup ideas	0.36	Kogan *et al.* (2017) and Section 2

B. Internally Calibrated Parameters

Parameter	*Description*	*Value*
λ	Innovation step size	1.040
ξ_S	Startup creation cost	0.183
κ_S	Development cost scale for startups	1.055
ξ_I	Research cost scale for incumbents	0.009
κ_I	Development cost scale for incumbents	0.338
χ	Search cost scale for incumbents	0.296
φ	Search cost curvature for incumbents	2.000

C. Model Fit

Targeted moment	*Model*	*Data*	*Data source*	*Identifies*
Growth rate	2.01%	2.00%	Jones (2016)	λ
Startup rate	25.4%	27.2%	Section 2	ξ_S
Avg. prob to impl. a startup idea (incumbents)	37.1%	36.6%	BDS, Section 2	κ_I
Avg. prob to impl. a startup idea (startups)	37.9%	36.6%	BDS, Section 2	κ_S
Growth contribution of startup ideas	28.0%	29.2%	Akcigit and Kerr (2018)	ξ_I
Percentage of startups that are acquired	5.5%	5.5%	Section 2	χ
Relative size of acquiring firms	1.5	2.8	Section 2	φ

startup rate of 27.2% implies that $\frac{9.4\%}{27.2\%} = 34.6\%$ of startups enter.[15] Thus, conditional on not being acquired, the average implementation probability of startups for their own ideas is $\frac{34.6\%}{100\%-5.5\%} = 36.6\%$. Furthermore, our analysis in Section 2 has shown that acquisitions do not appear to lead, on average, to a change in implementation rates. Therefore, we also target a 36.6% average implementation probability for startup ideas acquired by incumbents.[16] Our calibration implies that the implementation cost of incumbents is about 70% lower than the one of startups. Indeed, with equal implementation costs, our model would suggest that incumbents are considerably less likely to implement ideas than startups, due to the replacement effect. To account for the fact that startups and incumbents appear to implement at roughly the same rates in the data, our model assigns a large cost advantage to incumbents.

Fifth, we target the share of growth that can be attributed to startup ideas. To compute this number in the data, we rely on Akcigit and Kerr (2018), who find that innovation by entrants accounts for about 25% of economic growth.[17] This does not include startup ideas implemented by incumbents, however. In our data, the frequency of acquisitions is 1.6% (the product of a 27.2% startup rate and 5.5% of startups being acquired). Thus, the ratio of acquired startups to entering startups is $\frac{1.6\%}{9.4\%} \approx 0.17$. As the ideas of both groups of firms are of the same average quality, and are equally likely to be implemented, the total contribution of startup ideas to growth is $(1 + 0.17) \cdot 25\% \approx 29.2\%$. Though this moment is affected by a variety of parameters, it is particularly sensitive to the research costs of incumbents, ξ_I, which shifts the contribution of incumbents' own innovation to growth.

Finally, we target the size difference (in terms of sales) between acquiring firms and non-acquiring firms. Our data showed that acquirers were about 2.8 times larger than non-acquirers, suggesting selection effects in acquisition activity (for a related point, see David, 2020). This identifies the model parameter φ, the curvature in the search cost function, which governs how steeply costs increase for firms that search harder for startups.

[15] We obtain the 9.4% figure for the entry rate by taking the ratio of the number of new firms to the total number of firms, on average for the period 1981-2018, which we obtain from the Business Dynamics Statistics (BDS) data of the U.S. Census Bureau, available at https://www.census.gov/programs-surveys/bds/data.html.

[16] In the model, we compute the average probability of implementation conditional on an acquisition (I_A) and the average probability of implementation conditional on no acquisition (I_{NA}) as follows:

$$I_A \equiv \frac{\sum_\omega \sum_n m(\omega) \cdot m(n) \cdot \left(1 - \tilde{s}(\omega, n)\right) \cdot i_S(\omega)}{\sum_\omega \sum_n m(\omega) \cdot m(n) \cdot \left(1 - \tilde{s}(\omega, n)\right)}$$
$$I_{NA} \equiv \frac{\sum_\omega \sum_n m(\omega) \cdot m(n) \cdot \tilde{s}(\omega, n) \cdot i_A(\omega, n)}{\sum_\omega \sum_n m(\omega) \cdot m(n) \cdot \tilde{s}(\omega, n).}$$

[17] Garcia-Macia, Hsieh and Klenow (2019) find similar numbers.

4.2 Equilibrium properties

Before turning to the quantitative analysis, we discuss some key properties of the BGP equilibrium in this section (using the calibrated set of parameters).

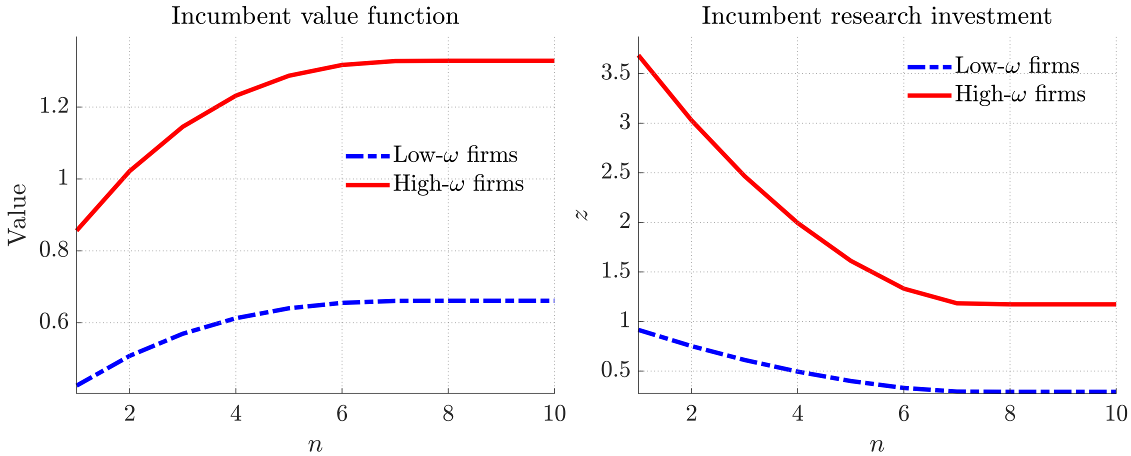

Figure 6: Value functions and research policy functions of incumbent firms, by firm type.

Figure 6 plots the value function v and the research policy function z for an incumbent firm. Firm value is increasing in quality ω and in the technology gap n. Moreover, firm value is concave in n, as the marginal effect of higher technology gaps on markups and profits gets smaller when the incumbent gets further ahead of its follower. Accordingly, once the incumbent is far ahead enough to charge the unconstrained monopoly markup, firm value no longer depends on the technology gap. The research investment of the firm, in turn, depends on the increments of the value function. Therefore, it is increasing in quality ω, and decreasing in the technology gap n. Note, however, that a firm which has reached the unconstrained monopoly markup still continues to invest into research: even though its markup cannot be increased further, the firm can still increase its market share by increasing productivity.

Figure 7 plots the acquisition surplus σ and the incumbent meeting probabilities s. In our model, acquisitions may have a positive surplus for two reasons. First, acquisitions could transfer an idea to a more efficient user (if $\kappa_I < \kappa_S$). Second, they allow the technology gap n to remain at least at its current value, instead of being potentially lowered through entry. The first motive reflects a socially useful transfer of ideas, while the second motive just preserves the rents of the incumbent firm (transferring part of them to the startup). A higher product quality ω and a higher technology gap n both imply greater benefits of transferring an idea to a better user and greater rents of maintaining the incumbent's position. Thus, the acquisition surplus is increasing in both variables, and firms with higher quality and higher technology gaps invest more resources into startup search.

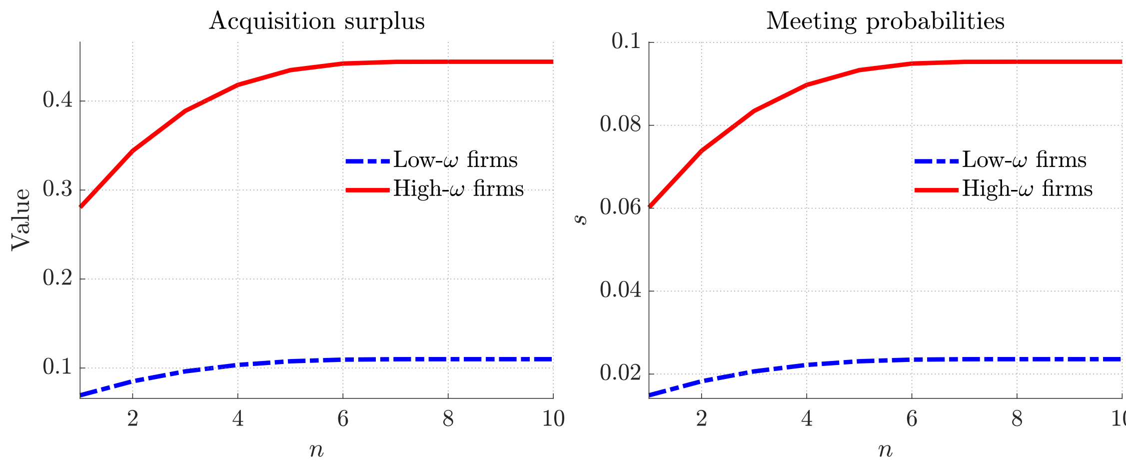

Figure 7: Acquisition surplus and meeting probabilities, by firm type.

Finally, the left panel of Figure 8 plots the implementation probabilities for a startup idea, distinguishing between the case in which the startup is not acquired and invests into implementation itself (i_S), and the case in which the startup is acquired and the incumbent invests into implementation (i_A). As shown in the previous section, incumbents face lower implementation costs than startups. Accordingly, at low levels of the technology gap, incumbents are more likely to implement a startup idea than the startup itself. As the technology gap increases, however, the marginal benefit of innovation for incumbents decreases (as the replacement effect becomes stronger). As a consequence, the implementation probability of ideas for incumbents falls below that of startups, and some acquisitions become killer acquisitions. On balance, acquisitions have a roughly neutral effect on the implementation probability, as imposed by our calibration.

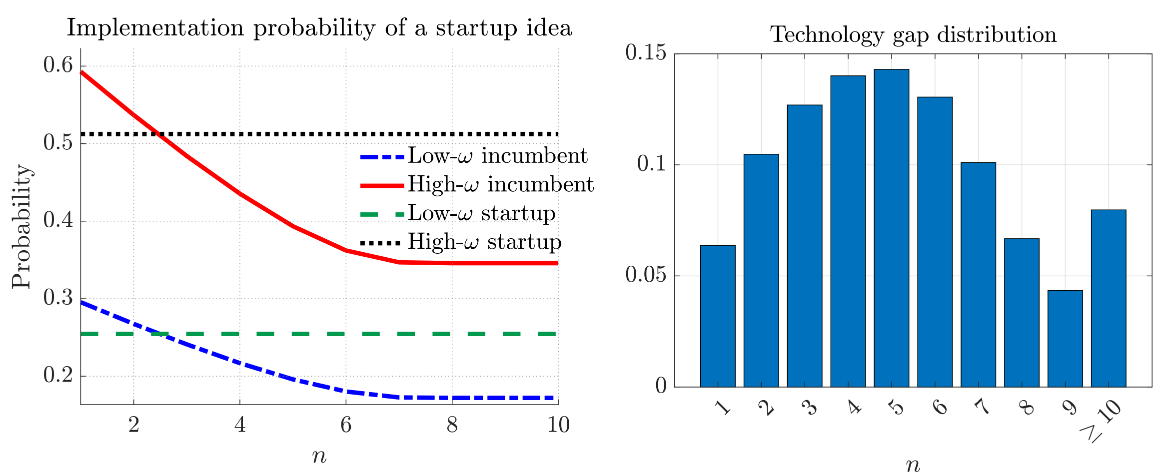

Figure 8: Development probabilities by firm type, and the invariant distribution of technology gaps.

The previous discussion shows that incumbent firm decisions about research, implementation and startup search crucially depend on the technology gap n. Therefore, the distribution of technology gaps across industries, shown in the right panel of Figure 7, is a crucial equilibrium object. This distribution is endogenous, shaped by the innovation choices of incumbents and startups. As we will see in the next sections, prohibiting or encouraging acquisitions will trigger shifts in this distribution.

4.3 The Aggregate Effects of Acquisitions

A useful decomposition Before discussing aggregate outcomes, it is worth investigating the sources of aggregate growth in somewhat greater detail. Our model admits a simple decomposition for changes in the growth rate between different BGPs.

$$\begin{aligned}\frac{g}{g^*} =& \sigma_I^* \cdot \frac{\text{Incumbent own innovation}}{\text{Incumbent own innovation}^*} + \\ & (1-\sigma_I^*) \cdot \left(\frac{\text{Startup rate}}{\text{Startup rate}^*} \cdot \frac{\text{Perc. of impl. startup ideas}}{\text{Perc. of impl. startup ideas}^*} \right),\end{aligned} \tag{31}$$

where x* stands for the baseline BGP value of variable x and σ_I^* stands for the BGP share of growth accounted for by incumbents' own innovation. Formally, the variables in this decomposition are given by

$$\begin{aligned}\text{Incumbent own innovation} =& \sum_{\omega,n} m(\omega,n) \cdot b_I(\omega,n) \\ \text{Startup rate} =& x \\ \text{Perc. of impl. startup ideas} =& \sum_{\omega,n} m(\omega,n) \cdot \Big(\widetilde{s}(\omega,n) \cdot i_A(\omega,n) + \Big(1-\widetilde{s}(\omega,n)\Big) \cdot i_S(\omega) \Big)\end{aligned}$$

Equation (31) shows that any change in the growth rate with respect to its baseline BGP value can be decomposed into three elements: changes in incumbent's own innovation behavior, changes in the startup rate, and changes in the percentage of startup ideas that are implemented. The weights in this expression are given by the baseline BGP share of growth accounted for by incumbents' own innovation, which is pinned down by our calibration target. When considering changes in growth between BGPs, we will keep referring to this decomposition.

Comparative statics To study the aggregate effect of acquisitions, we first consider our model's implication for changes in the search cost for startups χ. That is, we solve for the BGP equilibrium for different values of χ, keeping all other parameters at their baseline value. Recall that χ represents the level of frictions in the search for startups: a low value of this parameter implies low frictions and frequent acquisitions, while a high level implies high frictions and infrequent acquisition. Accordingly, as shown in the left panel of Figure 9, the equilibrium frequency of acquisitions is monotonically decreasing in χ.

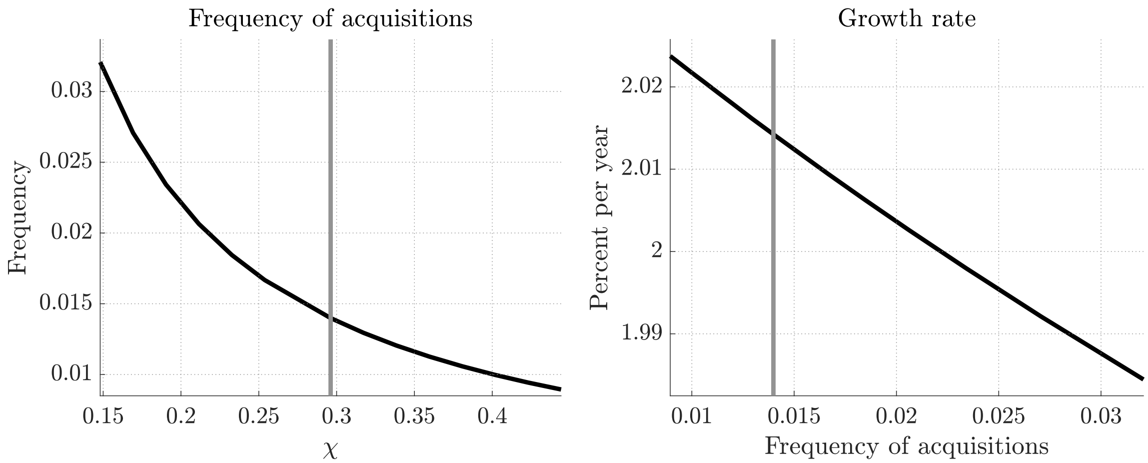

Figure 9: BGP equilibria for different values of search costs χ. The calibrated value is $\chi = 0.296$ (see Table 2), marked with a vertical line in the left plot. The right plot shows the reduced-form relationship between the frequency of acquisitions and growth. Again, the vertical line marks the baseline frequency of acquisitions.

The right panel of Figure 9 illustrates the main result of our paper: an increase in the frequency of acquisitions (due to lower search costs) implies a decrease in the growth rate. Note that to improve the readability of this and the subsequent plots, we directly show the frequency of acquisitions on the x-axis.

Figure 10 investigates the sources of this negative relationship. The top left panel plots the three sources of growth shown in the decomposition in equation (31), normalized to 1 at their baseline BGP level. It shows that when acquisitions increase, the arrival rate of startup ideas increases substantially: a doubling of the frequency of acquisitions implies roughly a 6% increase in this rate. Thus, acquisitions do have a strong incentive effect on startup creation in our model, which all else equal would imply that they are growth-enhancing. However, this positive effect is more than compensated by a decrease in incumbent's own innovation and in the percentage of startup ideas being implemented.

There are two main reasons for the change in these variables. First, higher acquisitions trigger a composition effect. As the percentage of acquired startups increases more strongly than the startup rate, the entry rate falls in our calibrated model. Creative destruction slows down, and the distribution of technology gaps shifts to the right (as shown in the bottom left panel of Figure 10). At a higher technology gap, the average incumbent has now less incentives to invest into research, or to implement its own and startup ideas. Second, a higher startup rate reduces the value for incumbents: even though incumbents can buy out startups and thereby avoid displacement, every acquisition implies a costly sharing of rents with the threatening startup. This decrease in incumbent value contributes to the fall in incumbent innovation, and explains why both incumbents and startups invest less into the implementation of ideas (as shown in the bottom right panel of Figure 10).

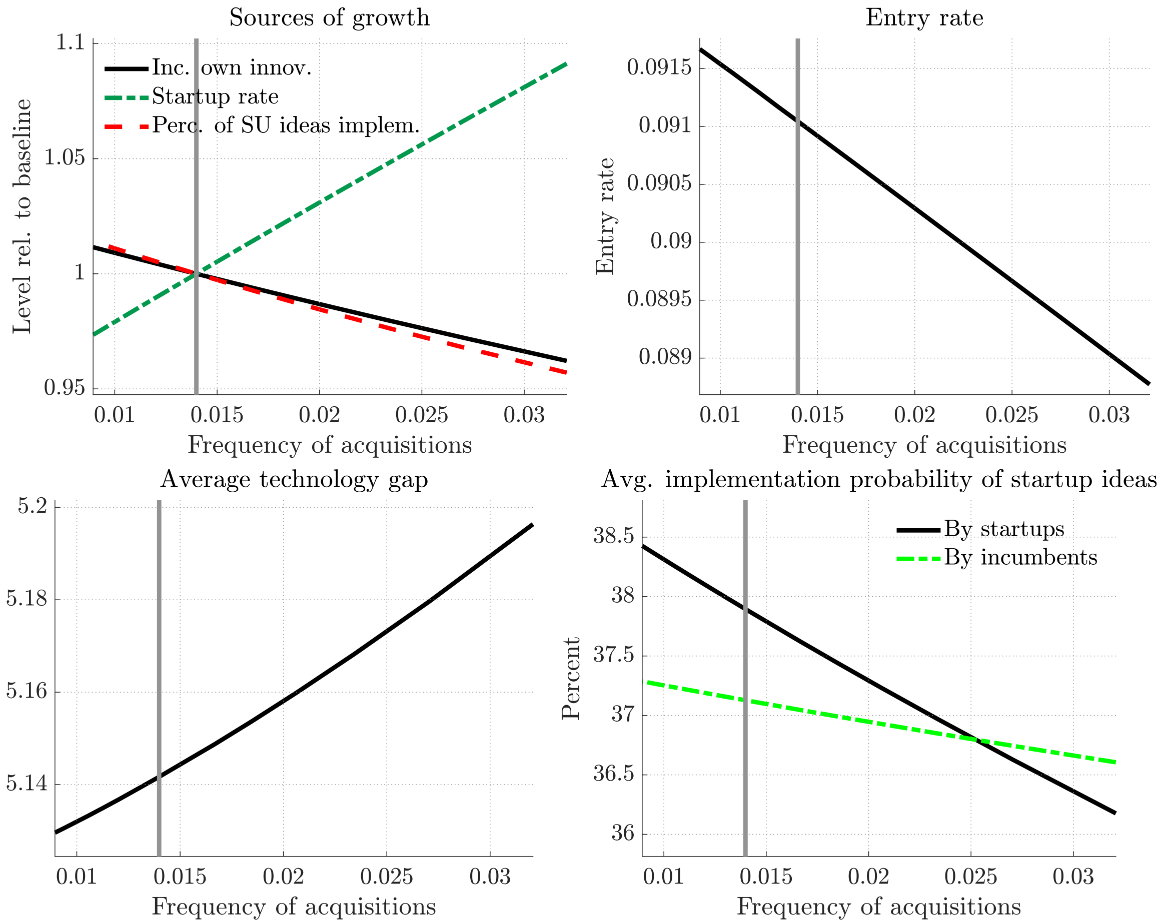

Figure 10: Important equilibrium outcomes for different values of search costs χ. All plots show the frequency of acquisitions on the x-axis (and all variation in this frequency is driven by changes in search costs). The vertical line marks the baseline frequency of acquisitions.

4.4 Policy

The comparative statics results shown in the previous section hint at a positive effect of stricter antitrust policy on growth. To confirm this impression, we now consider an actual policy experiment, by studying the effect of an acquisition ban.

Table 3 shows that this policy would lead to a slight increase in the aggregate growth rate, by 0.03 percentage points (or 1.4%) per year. In line with the intuitions developed above, this is the net effect of a 7.5% decrease in the arrival rate of startup ideas and an (overcompensating) 3.3% increase in the own innovation effort of incumbents and a 4.2% increase in the percentage of implemented startup ideas. Banning acquisitions also increases the entry rate and lowers the aggregate markup.

As discussed earlier, there are two reasons for which incumbents' own innovation and the implementation of startup ideas increase after the acquisition ban. First, with greater entry, the technology gap distribution shifts to the left (towards more innovative firms). Second, the disappearance of costly acquisitions increases the value of incumbents. To assess the relative strength of these two channels, we consider a counterfactual accounting exercise in which we only shift the distribution of technology gaps to its post-policy level, but keep all other variables at their baseline levels. This triggers a 0.007 percentage point increase in growth. Thus, roughly speaking, the shift in the distribution accounts for one quarter of the total growth effect, with the remainder due to changes in incumbent value.

Table 3: The effects of antitrust policy.

Outcome	**Baseline**	**Acq. Ban**	**% Change**
Growth rate	2.01%	2.04%	+1.4%
Incumbent own inn. rate	0.35	0.37	+3.3%
Startup rate	0.25	0.24	−7.5%
Percentage of imp. startup ideas	37.9%	39.4%	+4.2%
Entry rate	0.091	0.093	+1.9%
Share acquired startups	5.5%	0%	−100%
Aggregate markup	18.0%	17.9%	−0.5%

Overall, our results in this section suggest that the positive and negative effects of acquisitions on growth are roughly balanced (even though the negative effects do appear to be somewhat stronger). In the remainder of the paper, we explore the robustness of these baseline results to different choices for the targeted data moments and the externally calibrated parameters. This shows that in industries with more frequent acquisitions or lower competition, the negative effects of incumbents could be substantially larger than in this baseline.

4.5 Robustness

4.5.1 The role of calibration targets

First, we examine the role of different calibration targets for our quantitative conclusions. To do so, we will recalibrate our model by changing one calibration target at the time, leaving all other targets (and all externally set parameters) unchanged.

Figure 11 plots the results for this exercise for different values of the percentage of startups being acquired (which is directly linked to the frequency of acquisitions). The baseline value of this target was 5.5%, and each point on the x-axis corresponds to an alternative target. As the left-hand side panel shows, the effect of an acquisition ban is increasing in the frequency of acquisitions. With our baseline target, banning startup acquisitions increases growth by 1.4%. When acquisitions are twice as frequent, instead, banning acquisitions increases growth by about 3%. Intuitively, the more prevalent acquisitions are, the larger is their aggregate effect. Our model suggests that the elasticity of the growth effect to this target is roughly one.

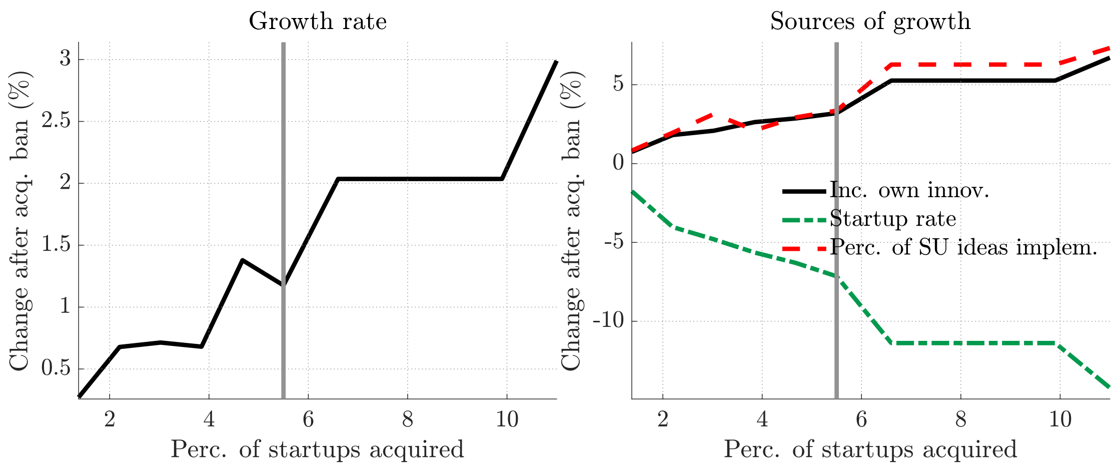

Figure 11: The role of calibration targets. This figure shows outcomes for different targets for the share of startups that are acquired (5.5% in the baseline, as marked by the vertical line). All other calibration targets and external parameter values are unchanged throughout.

Figure 12 instead examines the role of differences in implementation probabilities. In our baseline calibration, we targeted an implementation probability of startup ideas of 36.6% for both startups and incumbents. Here, we instead consider different targets for the average implementation probability of incumbents, keeping the one of startups unchanged. Thus, for any value lower than the baseline, the average acquisition is a killer acquisition, while for every value higher than the baseline, the average acquisition increases the probability that an idea gets implemented.

This figure illustrates that the prevalence of killer acquisitions is a key driver of our quantitative results. Indeed, targeting more killer acquisitions increases the negative effects of acquisitions on growth. For instance, when the average implementation probability of incumbents is around 15% (21 percentage points lower than the one of startups), banning startup acquisitions increases growth by almost twice as much as in the baseline. On the other hand, when the average implementation probability of incumbents is above 65%

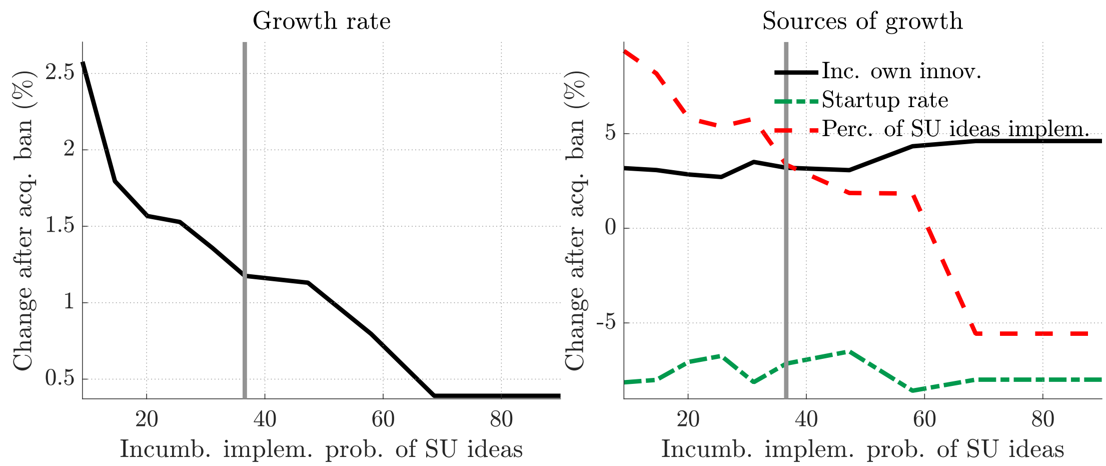

Figure 12: The role of calibration targets. This figure shows outcomes for different targets for the average probability that an incumbent implements a startup idea (36.6% in the baseline, as marked by the vertical line). All other calibration targets and external parameter values are unchanged throughout.

(more than 39 percentage points higher than the one of startups), an acquisition ban has virtually no effect on the growth rate. As Figure 12 shows, much of this variation is explained by the change in the percentage of implemented startup ideas: the more killer acquisitions there are, the more an acquisition ban increases this probability.

4.5.2 The role of external parameters

Figure 13 shows how our results depend on a key external parameter, the elasticity of substitution ε. To generate this figure, we solve for our model's BGP with different values for the parameter ε (leaving all other parameters unchanged).

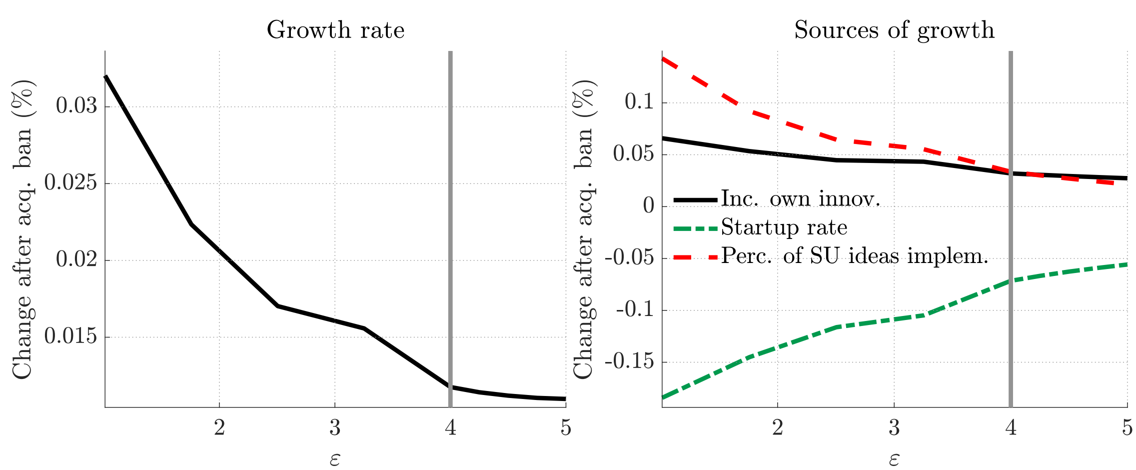

Figure 13: Robustness: the role of ε. This figure shows outcomes for different elasticities ε (4 in the baseline, as marked by the vertical line). All other parameter values are unchanged throughout.

As shown in the left panel of Figure 13, an acquisition ban has a stronger effect on growth when the elasticity of substitution is low. Indeed, with a low elasticity of substitution, firms care more about their relative productivity (the technology gap n) and not so much about their absolute productivity. In the extreme case of Cobb-Douglas aggregation (when ε tends to 1), absolute productivity is irrelevant, as firm sales do not depend on it. Thus, with a low elasticity of substitution, the elasticity of innovation spending to the technology gap is higher, and the leftward shift in technology gaps after an acquisition ban has a larger effect. Interpreting this result at the industry-level, it suggests that acquisitions are more problematic in industries in which products are less substitutable.

Overall, the robustness checks in this section imply that the negative effects of acquisitions are stronger when acquisitions are more frequent, when incumbents do not have large advantages in development, and when competition between incumbents is weak. Given the large variation in these characteristics across industries, one would therefore expect antitrust policy to have much stronger effects in some industries than in others.

5 Conclusion

In this paper, we assess the effect of startup acquisitions on productivity growth, using a macroeconomic model that takes into account positive effects (on the startup rate and idea transfers) and negative effects (killer acquisitions and spillovers on incumbents' own innovation incentives). We calibrate the model using micro-level data, and find that higher acquisitions increase the startup rate, by providing additional incentives for startup creation. However, this is more than compensated by a decrease in incumbent's own innovation and in the implementation probability of ideas. Accordingly, a policy that bans all startup acquisitions would increase the rate of growth by around 0.03 percentage points per year. The exact effect of the policy depends on the frequency of acquisitions, the prevalence of killer acquisitions and the degree of competition among incumbents.

References

AGHION, P., AKCIGIT, U. and HOWITT, P. (2014). What Do We Learn From Schumpeterian Growth Theory? In P. Aghion and S. N. Durlauf (eds.), *Handbook of Economic Growth*, *Handbook of Economic Growth*, vol. 2, Elsevier, pp. 515–563.

—, BERGEAUD, A., BOPPART, T., KLENOW, P. J. and LI, H. (2021). A Theory of Falling Growth and Rising Rents. *Mimeo*.

—, BLOOM, N., BLUNDELL, R., GRIFFITH, R. and HOWITT, P. (2005). Competition and Innovation: An Inverted-U Relationship. *The Quarterly Journal of Economics*, **120** (2), 701–728.

AKCIGIT, U., CELIK, M. A. and GREENWOOD, J. (2016). Buy, Keep, or Sell: Economic Growth and the Market for Ideas. *Econometrica*, **84** (3), 943–984.

— and KERR, W. R. (2018). Growth through Heterogeneous Innovations. *Journal of Political Economy*, **126** (4), 1374–1443.

ANDERSSON, M. and XIAO, J. (2016). Acquisitions of start-ups by incumbent businesses: A market selection process of "high-quality" entrants? *Research Policy*, **45** (1), 272 – 290.

ARROW, K. (1962). Economic Welfare and the Allocation of Resources for Invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, National Bureau of Economic Research, Inc, pp. 609–626.

BENA, J. and LI, K. (2014). Corporate innovations and mergers and acquisitions. *The Journal of Finance*, **69** (5), 1923–1960.

BESSEN, J. (2009). NBER PDP Project User Documentation: Matching Patent Data to Compustat Firms.

BLONIGEN, B. A. and PIERCE, J. R. (2016). *Evidence for the Effects of Mergers on Market Power and Efficiency*. NBER Working Papers 22750, National Bureau of Economic Research, Inc.

BOURREAU, M., JULLIEN, B. and LEFOUILI, Y. (2018). *Mergers and Demand-Enhancing Innovation*. TSE Working Papers 18-907, Toulouse School of Economics (TSE).

BRYAN, K. and HOVENKAMP, E. (2020). Antitrust limits on startup acquisitions. *Review of Industrial Organization*, **56** (4), 615–636.

CABRAL, L. M. B. (2018). *Standing on the Shoulders of Dwarfs: Dominant Firms and Innovation Incentives*. CEPR Discussion Papers 13115, C.E.P.R. Discussion Papers.

CALLANDER, S. and MATOUSCHEK, N. (2020). The Novelty of Innovation: Competition, Disruption, and Antitrust Policy. *Working paper*.

CAVENAILE, L., CELIK, M. A. and TIAN, X. (2020). The Dynamic Effects of Antitrust Policy on Growth and Welfare. *Working paper*.

CELIK, M. A., TIAN, X. and WANG, W. (2020). *Acquiring Innovation under Information Frictions*. Working paper, University of Toronto.

CUNNINGHAM, C., EDERER, F. and MA, S. (2020). Killer acquisitions. *Journal of Political Economy, forthcoming*.

DARUICH, D. (2020). The Macroeconomic Consequences of Early Childhood Development Policies. *Mimeo*.

DAVID, J. (2020). The Aggregate Implications of Mergers and Acquisitions. *Review of Economic Studies, forthcoming*.

DENICOLÒ, V. and POLO, M. (2021). Acquisitions, innovation and the entrenchment of monopoly. *Mimeo*.

DIMOPOULOS, T. and SACCHETTO, S. (2017). Merger activity in industry equilibrium. *Journal of Financial Economics*, **126** (1), 200 – 226.

FEDERICO, G., LANGUS, G. and VALLETTI, T. (2017). A simple model of mergers and innovation. *Economics Letters*, **157**, 136 – 140.

FUMAGALLI, C., MOTTA, M. and TARANTINO, E. (2020). *Shelving or developing? The acquisition of potential competitors under financial constraints*. Economics Working Papers 1735, Department of Economics and Business, Universitat Pompeu Fabra.

GALÍ, J. and MONACELLI, T. (2016). Understanding the gains from wage flexibility: The exchange rate connection. *American Economic Review*, **106** (12), 3829–68.

GARCIA-MACIA, D., HSIEH, C.-T. and KLENOW, P. J. (2019). How Destructive Is Innovation? *Econometrica*, **87** (5), 1507–1541.

GAUTIER, A. and LAMESCH, J. (2020). Mergers in the digital economy. *Information Economics and Policy*, p. 100890.

HAUCAP, J., RASCH, A. and STIEBALE, J. (2019). How mergers affect innovation: Theory and evidence. *International Journal of Industrial Organization*, **63**, 283 – 325.

JONES, C. I. (2016). The facts of economic growth. *Handbook of Macroeconomics*, vol. 2, Elsevier, pp. 3–69.

JOVANOVIC, B. and ROUSSEAU, P. L. (2002). The Q-Theory of Mergers. *American Economic Review*, **92** (2), 198–204.

KAMEPALLI, S. K., RAJAN, R. and ZINGALES, L. (2020). *Kill Zone*. Working Paper 27146, National Bureau of Economic Research.

KIM, J. D. (2020). The Paradox of Startup Acquisitions: Eliminating Competition and Creating New Competitors. *Working paper*.

KLETTE, T. J. and KORTUM, S. (2004). Innovating firms and aggregate innovation. *Journal of political economy*, **112** (5), 986–1018.

KOGAN, L., PAPANIKOLAOU, D., SERU, A. and STOFFMAN, N. (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, **132** (2), 665–712.

LENTZ, R. and MORTENSEN, D. T. (2016). Optimal growth through product innovation. *Review of Economic Dynamics*, **19**, 4 – 19, Special Issue in Honor of Dale Mortensen.

LETINA, I., SCHMUTZLER, A. and SEIBEL, R. (2020). *Killer Acquisitions and Beyond: Policy Effects on Innovation Strategies*. CEPR Discussion Papers 15167, C.E.P.R. Discussion Papers.

PETERS, M. (2020). Heterogeneous mark-ups, growth and endogenous misallocation. *Econometrica*, **88** (5), 2037–2073.

PHILLIPS, G. M. and ZHDANOV, A. (2013). R&D and the Incentives from Merger and Acquisition Activity. *The Review of Financial Studies*, **26** (1), 34–78.

RHODES-KROPF, M. and ROBINSON, D. T. (2008). The Market for Mergers and the Boundaries of the Firm. *The Journal of Finance*, **63** (3), 1169–1211.

SERU, A. (2014). Firm boundaries matter: Evidence from conglomerates and R&D activity. *Journal of Financial Economics*, **111** (2), 381 – 405.

WOLLMANN, T. G. (2019). Stealth Consolidation: Evidence from an Amendment to the Hart-Scott-Rodino Act. *American Economic Review: Insights*, **1** (1), 77–94.

The Aggregate Effects of Acquisitions on Innovation and Economic Growth

by Christian Fons-Rosen, Pau Roldan-Blanco and Tom Schmitz

Appendix Materials

A Derivations and Proofs

A.1 Normalized value functions and free entry condition

Using our guess for the value function, we can rewrite Equations (10) to (14) as

$$v^{\text{NoMeet}}(\omega,n) = (1-i_S(\omega))\cdot v(\omega,n), \tag{A.1}$$

$$v_S^{\text{NoMeet}}(\omega,n) = \max_{i_S}\left\{ i_S\cdot\left(\sum_{n_S=1}^{+\infty}\theta(n_S)\cdot\lambda^{n_S(\varepsilon-1)}\cdot v(\omega,n_S)\right) - \kappa_S\cdot i_S^{\psi}\right\} \tag{A.2}$$

$$\begin{aligned}\widetilde{\sigma}(\omega,n) = \max_{i_A}\Bigg\{ & v(\omega,n) + i_A\cdot\left(\sum_{n_S=1}^{+\infty}\theta(n_S)\cdot\lambda^{n_S(\varepsilon-1)}\cdot v(\omega,n+n_S) - v(\omega,n)\right) \\ & - \kappa_I\cdot i_A^{\psi}\Bigg\} - v^{\text{NoMeet}}(\omega,n) - v_S^{\text{NoMeet}}(\omega,n)\end{aligned} \tag{A.3}$$

$$v^{\text{Meet}}(\omega,n) = v^{\text{NoMeet}}(\omega,n) + \alpha\cdot\widetilde{\sigma}(\omega,n) \tag{A.4}$$

$$v_S^{\text{Meet}}(\omega,n) = v_S^{\text{NoMeet}}(\omega,n) + (1-\alpha)\cdot\widetilde{\sigma}(\omega,n). \tag{A.5}$$

In all of these expressions, small letters denote values that are normalized by relative productivity and aggregate GDP (e.g., $v^{\text{NoMeet}}(\omega,n)\cdot\widetilde{a}^{\varepsilon-1}\cdot Y_t = V_t^{\text{NoMeet}}(\omega,n,a)$, and so on). Using these expressions, we can rewrite the value of a startup - the right-hand side of equation (15) - as

$$\begin{aligned}&\mathbb{E}_t\left[s(\omega,n)\cdot V_{S,t}^{\text{Meet}}(\omega,n,a) + \Big(1-s(\omega,n)\Big)\cdot V_{S,t}^{\text{NoMeet}}(\omega,n,a)\right] \\ &= \sum_{\omega\in\Omega}\sum_{n=1}^{+\infty}\sum_{a\in\mathbb{A}_t} m(\omega)\cdot m(n)\cdot m_t(a)\cdot\Big[s(\omega,n)\cdot v_S^{\text{Meet}}(\omega,n) \\ &\qquad + \Big(1-s(\omega,n)\Big)\cdot v_S^{\text{NoMeet}}(\omega,n)\Big]\cdot\widetilde{a}^{\varepsilon-1}\cdot Y_t\end{aligned}$$

where $\mathbb{A}_t$ stands for the set of all productivities at instant t (note that because productivity evolves on a ladder, this set is always countable), and we have used the fact that the distributions of quality, technology gaps and productivity are independent. This independence allows us to rewrite the value of a startup as

$$Y_t \cdot \sum_{\omega\in\Omega}\sum_{n=1}^{+\infty} m(\omega)\cdot m(n)\cdot\left[s(\omega,n)\cdot v_S^{\text{Meet}}(\omega,n)+\Big(1-s(\omega,n)\Big)\cdot v_S^{\text{NoMeet}}(\omega,n)\right]\cdot\left(\sum_{a\in\mathbb{A}_t} m_t(a)\cdot\widetilde{a}^{\varepsilon-1}\right).$$

Next, note that that by definition,

$$\sum_{a\in\mathbb{A}_t} m_t(a)\cdot\widetilde{a}^{\varepsilon-1}=\int_0^1\left(\frac{a_{jt}}{A_t}\right)^{\varepsilon-1}\mathrm{d}j=1.$$

Replacing this into the previous expression yields equation (23) in the main text.

A.2 The invariant distribution of technology gaps

To determine the invariant distribution over technology gaps n, we build an *intensity* (also known as *infinitesimal generator*) matrix. For a homogeneous continuous-time Markov chain z_t taking values in some discrete space $\{z_1, z_2, \ldots, z_S\} \in \mathbb{R}^S$, a generator matrix $\boldsymbol{M}_z$ is defined by:

$$\boldsymbol{M}_z \equiv \begin{pmatrix} -\sum_{j\neq 1}\lambda_{1j} & \lambda_{12} & \ldots & \lambda_{1S} \\ \lambda_{21} & -\sum_{j\neq 2}\lambda_{2j} & \ldots & \lambda_{2S} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{S1} & \lambda_{S2} & \ldots & -\sum_{j\neq S}\lambda_{Sj} \end{pmatrix} \tag{A.6}$$

where $\lambda_{ij} \geq 0$ is the intensity rate for a z_i-to-z_j transition. Note that the diagonal elements of $\boldsymbol{M}_z$ collect outflows, while the off-diagonal elements collect inflows. Thus, each row of an infinitesimal generator matrix must add up to zero.

To build this matrix in our model, we assume $n \in \{1, 2, \ldots, n_{max}\}$, i.e. that the technology gap is bounded above by $n_{max} < +\infty$.

We denote by $m_t(n)$ the share of firms in state n at time t. The law of motion of $m_t(n)$ can be written as follows:

$$\frac{\partial \vec{m}_t}{\partial t} = \boldsymbol{M}_n^{\top}\vec{m}_t \tag{A.7}$$

To find the invariant distribution, we impose $\frac{\partial \vec{m}_t}{\partial t} = \vec{0}$ in equation (A.7) and solve for the unique solution of the system of linear equations holding $\sum_n m(n,k) = 1$.

What are the transition rates? For any transition from n to $n+1$, with $n+1 < n_{max}$, the transition rate is

$$\sum_{\omega} m(\omega) \cdot \left[z(\omega,n) \cdot i_I(\omega,n) + x \cdot \left(\widetilde{s}(\omega,n) \cdot i_A(\omega,n) \cdot \theta(1) + (1-\widetilde{s}(\omega,n)) \cdot i_S(\omega) \cdot \theta(n+1) \right) \right].$$

Transitions occur because of incumbent ideas, 1-step startup ideas implemented by incumbents, and $(n+1)$-step startup ideas implemented by startups.

For transitions from $n_{max} - 1$ to n_{max}, the transition rate is

$$\sum_{\omega} m(\omega) \cdot \left[z(\omega, n_{max}-1) \cdot i_I(\omega, n_{max}-1) + \right.$$
$$\left. x \cdot \left(\widetilde{s}(\omega, n_{max}-1) \cdot i_A(\omega, n_{max}-1) + (1-\widetilde{s}(\omega, n_{max}-1)) \cdot i_S(\omega) \sum_{n_S = n_{max}}^{+\infty} \theta(n_S) \right) \right].$$

The intuition is the same as before, but now any startup idea implemented by an incumbent brings us into n_{max}, as well as any startup idea of quality n_{max} or larger.

Next, for transitions from n to $n+k$, with $k > 1$ and $n+k < n_{max}$, we have a transition rate

$$\sum_{\omega} m(\omega) \cdot x \cdot \left(\widetilde{s}(\omega,n) \cdot i_A(\omega,n) \cdot \theta(k) + (1-\widetilde{s}(\omega,n)) \cdot i_S(\omega) \cdot \theta(n+k) \right).$$

These transitions can only occur because of startup ideas allowing an incumbent to take k steps or a startup to take $n+k$ steps.

For transitions from n to n_{max}, with $n < n_{max} - 1$, we have

$$\sum_{\omega} m(\omega) \cdot x \cdot \left(\widetilde{s}(\omega,n) \cdot i_A(\omega,n) \sum_{n_S = n_{max}-n}^{+\infty} \theta(n_S) + (1-\widetilde{s}(\omega,n)) \cdot i_S(\omega) \sum_{n_S = n_{max}}^{+\infty} \theta(n_S) \right).$$

These transitions happen when startup ideas allow an incumbent to take $n_{max} - n$ steps or more, or a startup to take n_{max} steps or more.

Finally, for downward transitions, from n_1 to n_2 with $n_1 > n_2$, we have a transition rate

$$\sum_{\omega} m(\omega) \cdot \left(x \cdot (1-\widetilde{s}(\omega,n_1)) \cdot i_S(\omega) \cdot \theta(n_2) \right)$$

Downward transition happen only when there is entry, and the entering startup takes n_2 steps on the quality ladder.

A.3 Aggregate ratios

To obtain equation (24) in the main text, we use the CES price index. The price of the final consumption good holds

$$P_t = 1 = \left(\int_0^1 \omega_{jt} \cdot p_{jt}^{1-\varepsilon} \mathrm{d}j\right)^{\frac{1}{1-\varepsilon}}$$
$$= \sum_{\omega\in\Omega} \sum_{n=1}^{+\infty} \sum_{a\in\mathbb{A}_t} m(\omega)\cdot m(n)\cdot m_t(a)\left[\omega\cdot\mu(n)\cdot w_t^{1-\varepsilon}\cdot a^{\varepsilon-1}\right]$$

As markups are only dependent on n, we can use the fact that $\sum_{\omega\in\Omega}\omega m(\omega) = 1$ and $\sum_{a\in\mathbb{A}_t} m_t(a)\cdot a^{\varepsilon-1} = A_t^{\varepsilon-1}$ to obtain the equation in the main text.

To derive the labor market clearing condition, we first note that the labor demand of each individual firm is

$$l_{jt} = \omega_{jt}\cdot\left(\mu(n_{jt})\right)^{-\varepsilon}\cdot w_t^{-\varepsilon}\left(a_{jt}\right)^{\varepsilon-1}\cdot Y_t.$$

Integrating over all producers and imposing labour market clearing, we get

$$L = \left[\int_0^1 \omega_{jt}\cdot\left(\mu(n_{jt})\right)^{-\varepsilon}\cdot\left(a_{jt}\right)^{\varepsilon-1}\right]\cdot w_t^{-\varepsilon}\cdot Y_t$$
$$= \left[\sum_{\omega\in\Omega}\sum_{n=1}^{+\infty}\sum_{a\in\mathbb{A}_t} m(\omega)\cdot m(n)\cdot m_t(a)\cdot\left(\omega\cdot\left(\mu(n)\right)^{-\varepsilon}\cdot a^{\varepsilon-1}\right)\right]\cdot w_t^{-\varepsilon}\cdot Y_t$$
$$= \left[\sum_{n=1}^{+\infty} m(n)\cdot\left(\mu(n)\right)^{-\varepsilon}\right]\cdot A_t^{\varepsilon-1}\cdot w_t^{-\varepsilon}\cdot Y_t,$$

which immediately yields equation (25) in the main text.

A.4 Growth Rate

As shown in the main text, on the BGP, aggregate output, consumption and wages all grow at the same rate as average productivity A_t. To derive the growth rate of average productivity, we first note

$$\ln(A_t) = \frac{1}{\varepsilon-1}\cdot\ln\left(\int_0^1 a_{jt}^{\varepsilon-1}dj\right).$$

Now, consider an infinitesimally small time period dt. In this period, every product in state (ω, n) has a probability $b(\omega, n, k)\cdot dt$ of seeing its productivity increase by a factor λ^k, where $b(\omega, n, k)$ is defined as:

$$b(\omega, n, k) \equiv \begin{cases} b_I(\omega, n) + \theta(1) \cdot b_S(\omega, n) & \text{if } k = 1 \\ \theta(k) \cdot b_S(\omega, n) & \text{if } k \geq 2 \end{cases}$$

and (b_I, b_S) are the arrival rates of innovation by incumbents and startups defined in the main text. Therefore, applying the law of large numbers, we can write average productivity at instant $t + dt$ as

$$\begin{aligned} \ln(A_{t+dt}) = \frac{1}{\varepsilon - 1} \cdot \ln \Bigg[\sum_{a \in \mathbb{A}_t} \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \cdot m_t(a) \cdot \Bigg(\bigg(1 - \sum_{k=1}^{+\infty} dt \cdot b(\omega, n, k) \bigg) \cdot a^{\varepsilon - 1} \\ + \sum_{k=1}^{+\infty} dt \cdot b(\omega, n, k) \cdot \lambda^{k \cdot (\varepsilon - 1)} \cdot a^{\varepsilon - 1} \Bigg) \Bigg] \\ = \ln(A_t) + \frac{1}{\varepsilon - 1} \cdot \ln \Bigg[\sum_{a \in \mathbb{A}_t} \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \cdot m_t(a) \cdot \left(\frac{a}{A_t} \right)^{\varepsilon - 1} \\ + dt \cdot \sum_{a \in \mathbb{A}_t} \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \cdot m_t(a) \cdot \left(\frac{a}{A_t} \right)^{\varepsilon - 1} \sum_{k=1}^{+\infty} b(\omega, n, k) \cdot \left(\lambda^{k \cdot (\varepsilon - 1)} - 1 \right) \Bigg]. \end{aligned}$$

Using again that fact that relative productivity aggregates up to 1, and dividing by dt, we get

$$\frac{\ln(A_{t+dt}) - \ln(A_t)}{dt} = \frac{1}{\varepsilon - 1} \cdot \frac{\ln \left(1 + dt \cdot \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega) \cdot m(n) \sum_{k=1}^{+\infty} b(\omega, n, k) \cdot \left(\lambda^{k \cdot (\varepsilon - 1)} - 1 \right) \right)}{dt}.$$

Taking the limit as dt goes to 0 (and using the fact that $\lim_{x \to 0} \frac{\ln(1+x)}{x} = 1$), we get:

$$\frac{\dot{A}_t}{A_t} = \frac{1}{\varepsilon - 1} \cdot \left(\sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} \left[\sum_{k=1}^{+\infty} b(\omega, n, k) \cdot \left(\lambda^{k(\varepsilon - 1)} - 1 \right) \right] \right).$$

The term in square brackets is the overall innovation rate in state (ω, n), on average across step sizes. Note:

$$\sum_{k=1}^{+\infty} b(\omega, n, k) \cdot \left(\lambda^{k(\varepsilon - 1)} - 1 \right) = b_I(\omega, n) \cdot \left(\lambda^{\varepsilon - 1} - 1 \right) + b_S(\omega, n) \cdot \left[\sum_{k=1}^{+\infty} \theta(k) \cdot \left(\lambda^{k(\varepsilon - 1)} - 1 \right) \right].$$

Using that $\theta(k) = \frac{\gamma^{k-1}}{(k-1)!} \cdot e^{-\gamma}$ and $\sum_{k=1}^{+\infty} \theta(k) = 1$, we can write the term in brackets from the last expression as follows:

$$
\begin{aligned}
\sum_{k=1}^{+\infty} \theta(k) \cdot \left(\lambda^{k(\varepsilon-1)} - 1\right) &= e^{-\gamma} \cdot \left(\sum_{k=1}^{+\infty} \frac{\gamma^{k-1}}{(k-1)!} \cdot \lambda^{k\cdot(\varepsilon-1)}\right) - 1 \\
&= \lambda^{\varepsilon-1} \cdot e^{-\gamma} \cdot \left(\sum_{k=1}^{+\infty} \frac{\left(\gamma \cdot \lambda^{\varepsilon-1}\right)^{k-1}}{(k-1)!}\right) - 1 \\
&= \lambda^{\varepsilon-1} \cdot e^{\gamma\cdot(\lambda^{\varepsilon-1}-1)} - 1
\end{aligned}
$$

where, to go from the second to the third lines, we have used that the term in parenthesis on the second line is equal to $e^{\gamma\cdot\lambda^{\varepsilon-1}}$, because $e^{-\gamma\cdot\lambda^{\varepsilon-1}} \cdot \frac{\left(\gamma\cdot\lambda^{\varepsilon-1}\right)^{k-1}}{(k-1)!}$ equals the probability of a Poisson distribution with parameter $\gamma \cdot \lambda^{\varepsilon-1}$. Putting everything together, we finally obtain equation (30).

B Numerical Appendix

B.1 Solution Algorithm

To solve for the stationary solution, recall from equation (16) that the value of a firm in state (ω, n) depends on three aggregate variables: the wage-to-productivity ratio, A_t/w_t, the growth rate, g, and the arrival rate of startup ideas, x. As we argue in the text, all three of these variables are constant along the BGP. The solution algortihm therefore requires finding a fixed point for these three aggregates.

The following sketches the steps of our algorithm:

- Guess a value for the aggregate productivity-wage ratio, $\frac{A_t}{w_t}$.
- Guess a value for the growth rate g and the startup rate x.
- Given these guesses, solve the incumbent's problem using value function iteration.
- Compute the invariant distribution $m(n)$, the implied free entry condition and the implied growth rate. Use these to update the guesses for growth rate and arrival rate of startup ideas, and repeat until convergence.
- With the result, compute the productivity-to-wage ratio and update the guess.

B.2 Estimation Procedure and Global Identification

Next, we explain the estimation procedure and present a global identification test for the calibration exercise presented in Section 4.1.

We seek to find the set of M parameters, collected in the vector $\boldsymbol{\theta}$, that minimizes the distance between M moments generated by the model and their counterparts in the data. The distance function is:

$$\mathcal{D}(\boldsymbol{\theta}) \equiv \sum_{m=1}^{M} \frac{|\text{Moment}_m(\text{Model}, \boldsymbol{\theta}) - \text{Moment}_m(\text{Data})|}{0.5\,|\text{Moment}_m(\text{Model}, \boldsymbol{\theta})| + 0.5\,|\text{Moment}_m(\text{Data})|}$$

To perform such a minimization, rather than relying on gradient-based methods, we use an algorithm that efficiently searches over a large region of the parameter space and searches for the model solution that yields the lowest distance.

In particular, first we create a large M-dimensional hyper-cube $\mathcal{P}$ in the parameter space. Then, we pick quasi-random realizations from it using a Sobol sequence, which successively forms finer uniform partitions of the parameter space. For a sufficiently large number of Sobol draws, this routine efficiently and comprehensively explores every corner of $\mathcal{P}$. For each parameter evaluation, we then solve the model and store its results in a matrix. For this step, we use a high performance computer (HPC), allowing us to parallelize the procedure into hundreds of separate CPUs, thereby saving us an enormous amount of computation time. After N draws (in practice, $N \approx 1.5$ million), we have a $N \times M$ matrix $\boldsymbol{R}$ of results and a $N \times M$ matrix $\boldsymbol{\Theta} \in \mathcal{P}$ of the corresponding parameters. We then select the row vector $\widehat{\boldsymbol{\theta}} \in \boldsymbol{\Theta}$ for which $\mathcal{D}(\widehat{\boldsymbol{\theta}}) \leq \mathcal{D}(\boldsymbol{\theta})$, $\forall \boldsymbol{\theta} \neq \widehat{\boldsymbol{\theta}}$.

The advantage of this method over other estimation techniques is that the model-generated data contained in the $(\boldsymbol{R}, \boldsymbol{\Theta})$ matrices can be exploited to obtain information about identification. Particularly, we implement the following procedure, adapted from Daruich (2020). First, for each parameter p, we select a target moment m which we believe is particularly sensitive to the parameter. Note that, because of the Sobol routine, for each given value of p there is a distribution of values for m resulting from underlying random variation in all the remaining $M-1$ parameters. Using this fact, we then divide the support of p into 50 quantiles, and compute the 25th, 50th and 75th percentiles of this underlying distribution at each quantile.

We may now study how sensitive m is to changes in p by exploring the properties of how the moment's distribution behaves across different values for p. We say that p is well-identified by m when (i) the distribution changes across quantiles of p, (ii) the rate of this change is high, and (iii) the inter-quartile range of the m distribution is small throughout the support for p. Criterion (i) implies that m is *sensitive* to variation in p, (ii) gives an idea of *how strong* this relationship is, and (iii) implies that *other parameters* are relatively unimportant to explain it. Importantly, as all the remaining parameters are not fixed throughout this analysis but rather are varying in a random fashion, this method gives us a global view of identification and, therefore, presumably outperforms identification methods based on local elasticities (that is, based on moment pseudo-derivatives obtained by keeping the remaining parameters fixed at their calibrated values).

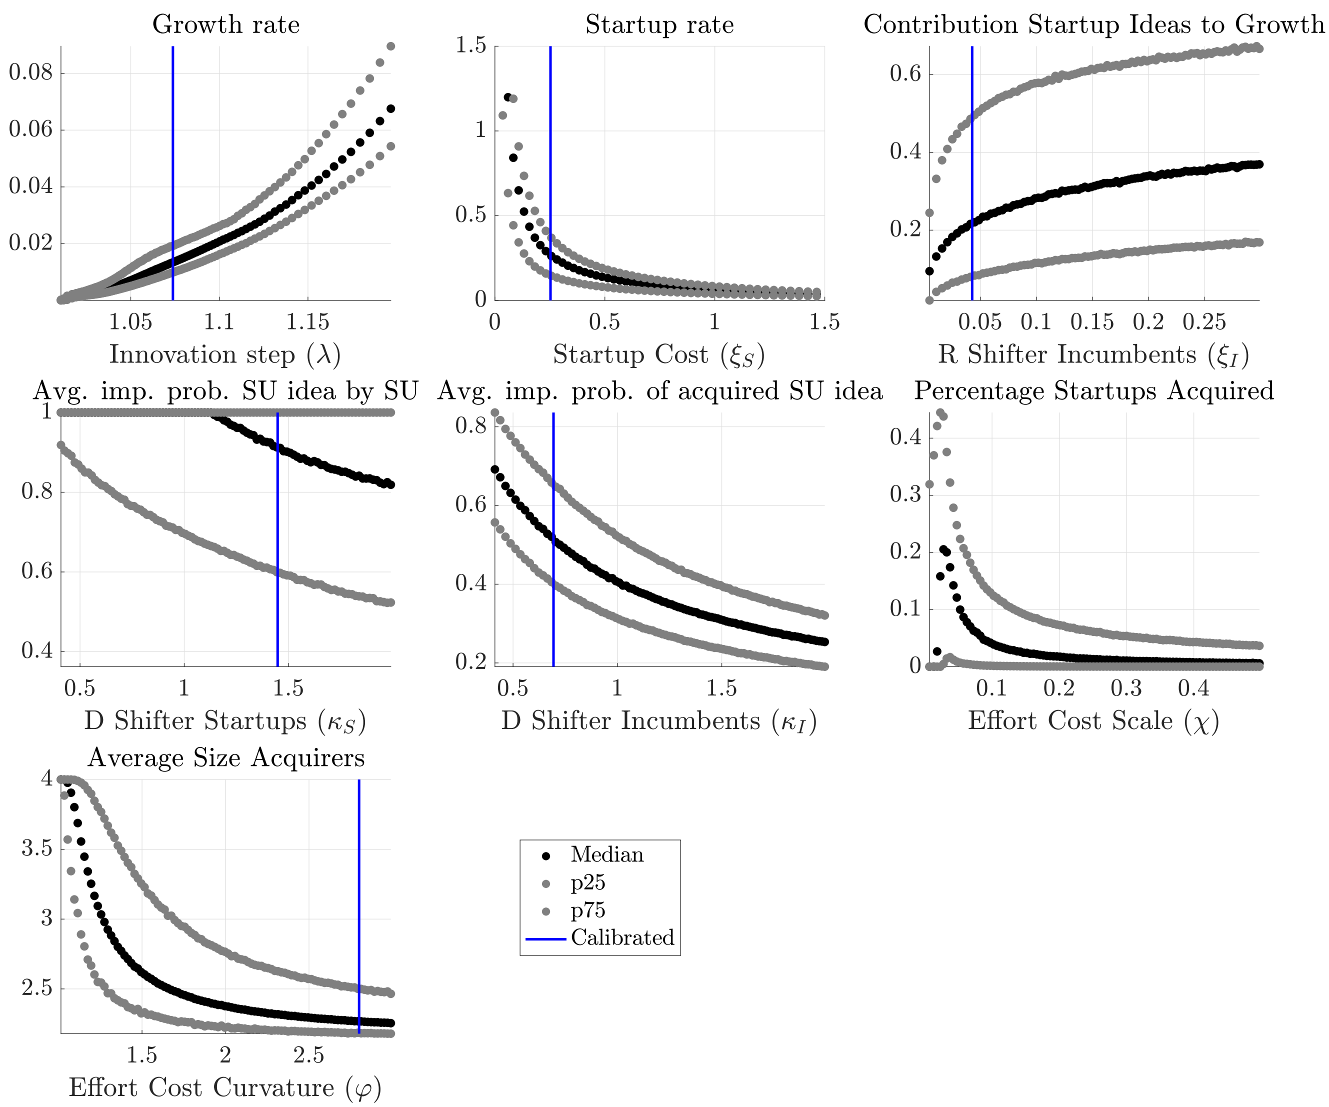

Figure B.1: Global identification results.

Figure B.1 presents the results from the global identification procedure explained above, where we have associated each targeted moment with the parameter that the moment most plausibly identifies (the same pairing as in Table 2 and in our verbal discussion in Section 4.1). All in all, we find that the parameters of the model are well-identified by criteria (i) and (ii) above and, with some exceptions, by criterion (iii).

C Additional Stylized Facts

Table C.1: Number of M&A transactions between 1980 and 2006

		Public Target		
		0	1	
Public	0	28,138	5,467	33,605
acquirer	1	24,398	4,253	28,651
		52,536	9,720	62,256

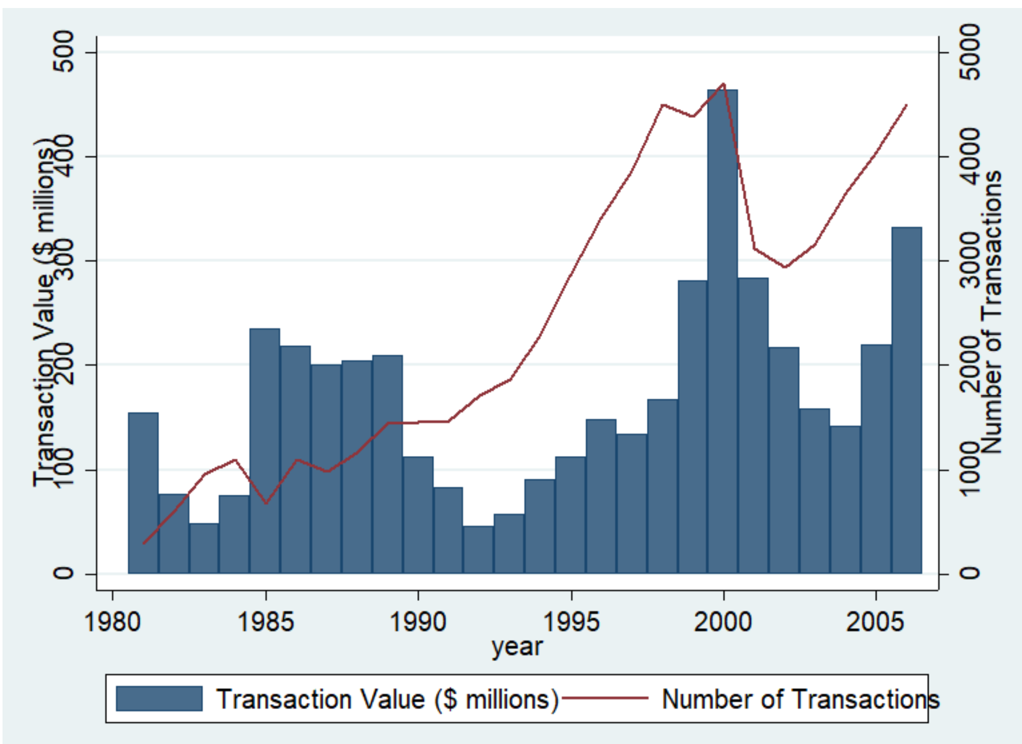

Figure C.1: Number of Transactions and Value per Year

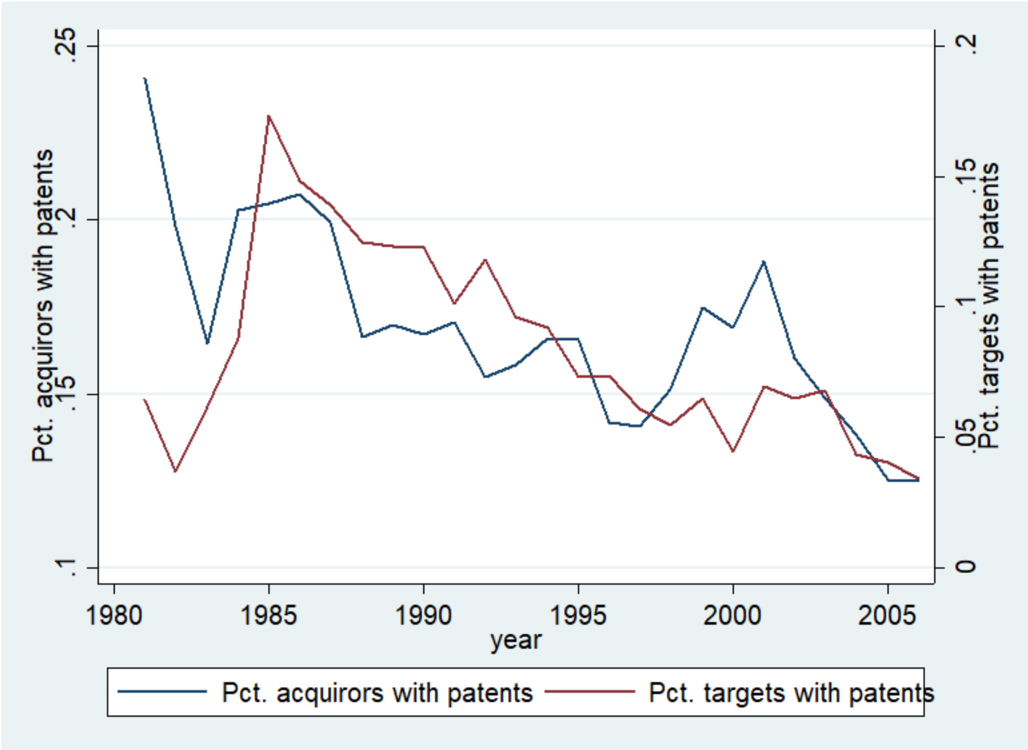

Figure C.2: Fraction of firms holding patents

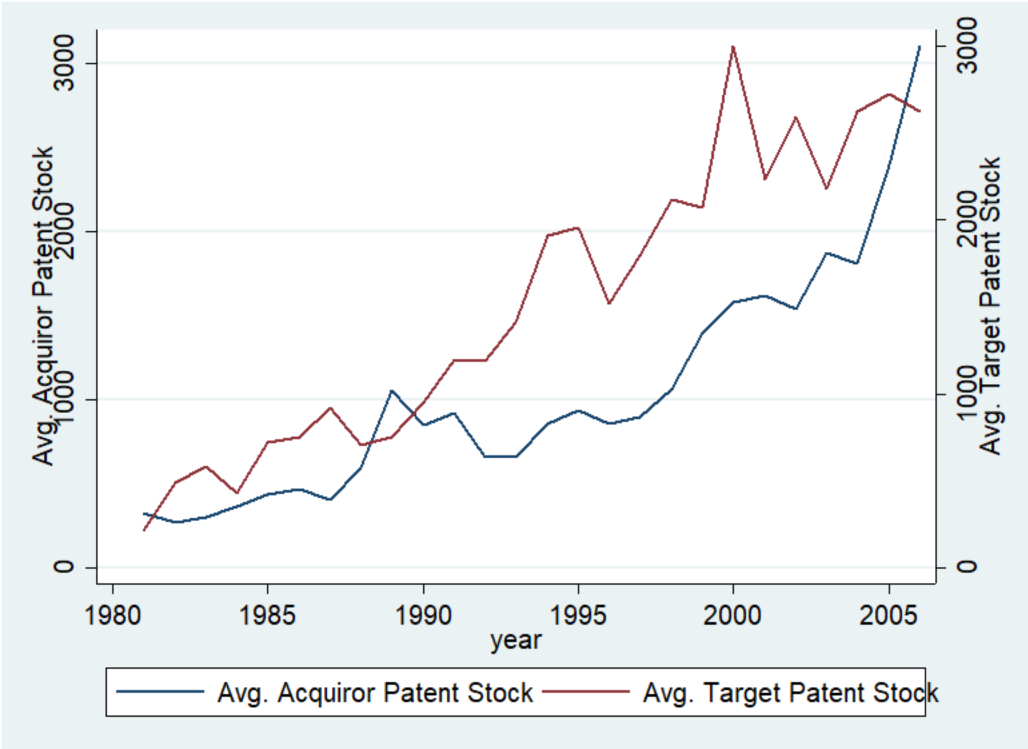

Figure C.3: Stock of patents conditional on holding patents