

The Effects of Startup Acquisitions on Innovation and Economic Growth*

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

First draft: February 2021; This version: January 9, 2023

Abstract

Innovative startups are frequently acquired by large incumbent firms. On the one hand, these acquisitions provide an incentive for startup creation and may transfer ideas to more efficient users. On the other hand, incumbents might acquire startups just to “kill” their ideas, and acquisitions can erode incumbents’ own innovation incentives. Our paper aims to assess the net effect of these forces. To do so, we build an endogenous growth model with heterogeneous firms and acquisitions, and calibrate its parameters by matching micro-level evidence on startup acquisitions and patenting in the United States. Our calibrated model implies that acquisitions raise the startup rate, but lower incumbents’ own innovation as well as the percentage of implemented startup ideas. The negative forces are slightly stronger. Therefore, a ban on startup acquisitions would increase growth by 0.03 percentage points per year, and raise welfare by 1.8%.

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

JEL Classification: O30, O41, E22.

*Paolo Mengano and Sutanuka Roy provided excellent research assistance. We thank Jorge Guzman for his help with the [Guzman and Stern \(2020\)](#) data, and Igor Letina, Huiyu Li and Ludovic Panon for helpful discussions. We also thank Gadi Barlevy, Vasco Carvalho, Laurent Cavenaile, Joel David, Maarten De Ridder, Florian Ederer, Chiara Fumagalli, Hugo Hopenhayn, Chad Jones, Danial Lashkari, Enrico Moretti, Jane Olmstead-Rumsey, Michael Peters, Joshua Weiss, Liangjie Wu, seminar participants at Banco de España, Bocconi, CEMFI, Chicago Fed, EIEF, Jinan, Kent, Sabanci, UAB, UB, UC Berkeley, UC Santa Cruz and UPF-CREi, and conference participants at SED 2021, EEA-ESEM 2021, EWMES 2021, the CEPR 4th Joint Research Conference on Firm Financing, Organization and Dynamics, T2M 2022, Bristol Macro Workshop 2022, Lisbon Macro Workshop 2022, BSE Summer Forum 2022, MMF 2022, the V Workshop of the Spanish Macroeconomic Network, and the 5th California Macro Conference 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 & Max Planck Institute for Innovation and Competition in Munich, 5200 North Lake Rd. Merced, CA 95343, United States, cfons-rosen@ucmerced.edu.

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

[§]Queen Mary University of London and CEPR, Mile End Road, London E1 4NS, United Kingdom, t.schmitz@qmul.ac.uk.

1 Introduction

Startups are a major source of innovative ideas. Therefore, they make a substantial contribution to aggregate productivity growth in the United States (see e.g. [Decker, Haltiwanger, Jarmin and Miranda, 2016](#)). However, many successful startups never grow into large independent firms, as they are acquired early on by older incumbents. For these incumbents, startup acquisitions are often a routine activity: several of the largest American firms have bought hundreds of startups over the last decade.¹

In recent years, regulators have viewed these operations with increasing skepticism. In 2020, the Federal Trade Commission (FTC) announced an inquiry into several high-profile startup acquisitions and filed lawsuits against two large incumbents, Google and Facebook.² In 2021, Congress introduced the bipartisan Platform Competition and Opportunity Act, which aims to prohibit acquisitions by certain large technology platforms, unless these firms can “*demonstrate that [they are] not acquiring a direct, nascent, or potential competitor, enhancing a market position, or enhancing [their] ability to maintain a market position*”.³

While such actions are often motivated by concerns about market power, regulators have also grown increasingly nervous about potential negative effects of startup acquisitions on innovation, both in the Tech sector and in the broader economy. A commonly used argument is that incumbents engage in “killer acquisitions” (a term coined by [Cunningham, Ederer and Ma, 2021](#)), acquiring startups only to kill innovative ideas that threaten the incumbent’s business. Industry lobbyists have pushed back and argue that acquisitions actually foster innovation because they motivate founders to create startups in the first place, and because incumbents are better prepared to commercialize startup ideas.

In this debate, the arguments of both sides have merit and are supported by some empirical evidence. To make informed decisions, policy makers therefore need to consider the balance of the positive and negative effects of startup acquisitions on innovation and productivity growth. To the best of our knowledge, no such comprehensive assessment is available to date. Our paper aims to fill this gap.

To do so, we develop an endogenous growth model with heterogeneous firms and startup acquisitions. Our model features the main forces discussed in the public debate, and also highlights the less commonly discussed (but no less important) effects of acquisitions

¹According to a report by the Federal Trade Commission (FTC), the Tech giants Amazon, Apple, Facebook, Google and Microsoft have acquired more than 600 small firms between 2010 and 2019 ([FTC, 2021](#)).

²A report on this inquiry was released in September 2021 ([FTC, 2021](#)). The FTC sued Google and Facebook in October and December 2020, asking Facebook to undo its acquisitions of Instagram and WhatsApp.

³See <https://itif.org/publications/2022/01/31/platform-competition-and-opportunity-act-solution-search-problem>.

on incumbents' own innovation. We discipline the model by calibrating it to micro data on acquisitions and patenting in the US. Our calibrated model implies that the negative forces slightly dominate, so that policies reducing acquisitions increase productivity growth.

Our model builds on the Schumpeterian growth framework. It features a continuum of incumbent firms, each producing a differentiated product and investing into innovation to increase its productivity gap over potential competitors. A large mass of non-producing startups, in turn, seeks to innovate to displace incumbents and enter. We introduce two new elements into this setting. First, we assume that innovation is a two-stage process: firms need to invest resources to invent new ideas, and then again to implement these ideas. This allows us to capture that incumbents might be more or less likely to implement an idea than the startup that came up with it. Second, we introduce startup acquisitions. In our baseline model, incumbents invest resources into searching for startups that threaten to displace them. When their search is successful, they can do an acquisition.

The model reflects the multiple channels through which startup acquisitions affect innovation and growth. Precisely, we show that any change in the growth rate from its baseline value (triggered, for instance, by a shock to search costs or a change in antitrust policy) can be decomposed into three margins: (i) changes in the startup rate (i.e., the number of new startups created in a given time period); (ii) changes in the sales-weighted percentage of startup ideas that are implemented; and (iii) changes in the own innovation rate of incumbents. The relative importance of these three margins depends on the initial shares of startup and incumbent ideas in aggregate productivity growth.

What drives changes in these three margins? First, in our model, shocks that increase the frequency of acquisitions generally also increase the startup rate. Indeed, startups only sell if the acquisition price exceeds their outside option of remaining independent. Thus, a higher frequency of acquisitions increases the value of creating a startup and leads to more startups being created in equilibrium. This positive effect mirrors the widespread conviction in the business world that acquisitions are a desirable outcome for startups.⁴

Second, the impact of acquisitions on the implementation of startup ideas is a priori unclear. On the one hand, incumbents might have lower implementation costs than startups.⁵ When this effect dominates, more frequent acquisitions lead, all else equal, to more implemented startup ideas. On the other hand, our model also features a “replacement

⁴In fact, numerous guides advise entrepreneurs how to position their startups in order to be acquired. 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>.

⁵Such a cost advantage might be due to economies of scale and scope, a larger customer base, a better access to capital markets, or greater business experience.

effect” as in [Arrow \(1962\)](#): an incumbent’s marginal benefit from implementing a startup idea is smaller than the one of the startup itself, providing a motive for “killing” an idea that would have been implemented by an independent startup. When this effect dominates, more frequent acquisitions might lead to fewer implemented startup ideas.

Finally, the own innovation rate of incumbents responds to startup acquisitions through several general equilibrium channels. A higher frequency of acquisitions attracts more startups, which means that incumbents are more frequently threatened with displacement. Even if incumbents manage to defuse this threat through an acquisition, this is costly, as they need to pay off the startup. The higher startup rate therefore lowers the value of incumbents and their incentives to innovate. Furthermore, incumbent innovation is also affected by a composition effect. In our model, incumbents with a high productivity gap over their potential competitors have low innovation incentives. The share of such firms in the economy crucially depends on entry (which resets productivity gaps), and entry is in turn affected by the frequency of acquisitions.

Our paper’s objective is to assess the quantitative importance of the three margins listed above. To do so, we calibrate our model to match micro data on innovation and acquisitions. Crucially, our targeted moments include direct causal evidence on the effect of acquisitions on the implementation of ideas.

We rely on two data sources. First, we use data collected by [Guzman and Stern \(2020\)](#), which covers the universe of startups created in 32 US states between 1988 and 2008. This data allows us, for instance, to compute the percentage of startups that are acquired. However, it does not identify the firms in the actual acquisition deals or follow their outcomes over time. Therefore, we construct a new data set by combining information on acquisitions (from the ThomsonONE M&A database), patents (from the NBER Patent Data Project) and accounting data (from Compustat).

Using this new data set, we study the impact of acquisitions on the implementation of ideas, one of the key margins in our model. To measure implementation, we track patent citations. We interpret an increase in the citations received by a startup patent after acquisition as evidence for the acquisition increasing the likelihood of implementation, and a decrease as evidence for the acquisition lowering this likelihood. To control for selection, we use a matching algorithm, assigning each treatment patent to a group of control patents with similar characteristics. For the average deal, we find that an acquisition increases citations of the acquired patent by around 22%, implying that incumbents have a comparative advantage for implementation that outweighs the Arrow replacement effect. In line with our model, we find that the boost to citations is lower if the acquirer has a high market share, or if the acquirer and the startup belong to the same industry.

We calibrate our model's parameters to match our regression results and other moments. To study the link between acquisitions and productivity growth, we then solve the model for different values of incumbents' search costs for startups. We find that high search costs (implying infrequent acquisitions) are associated with high growth, while low search costs (implying frequent acquisitions) are associated with low growth. Decomposing these differences into the three margins highlighted above, we find that more frequent acquisitions are associated with a higher startup rate, but also with lower rates of incumbent innovation and, surprisingly, a lower percentage of implemented startup ideas. The latter negative effects are due to a general equilibrium feedback: the higher startup rate erodes the value of incumbents, lowering innovation incentives for incumbents and non-acquired startups.⁶ As incumbent ideas are the main source of growth in our calibration (in line with the evidence in [Akcigit and Kerr, 2018](#) and [Garcia-Macia, Hsieh and Klenow, 2019](#)), the fall in incumbents' own innovation quantitatively dominates and drags the growth rate down.

In line with these results, we find that a ban on startup acquisitions would increase the aggregate growth rate by about 0.03 percentage points by year, and increase consumption-equivalent welfare by 1.8%. The ban lowers the startup rate, but this is more than compensated by an increase in incumbent innovation and in the percentage of implemented startup ideas. Partial bans have a similar, but somewhat smaller impact.

Finally, we explore the robustness of these results. Using our baseline model, we provide a range of potential results by considering different values for each of our calibration targets. This shows that higher values for the effect of acquisitions on idea implementation, a larger share of startup ideas in overall growth and higher bargaining power for incumbents would all somewhat dampen the negative effects of acquisitions. We also develop an extension of our model in which we allow incumbents to acquire startups that do not threaten to displace them. Results are quantitatively similar to the baseline, but underline that the negative effects are driven by acquisitions of competing startups.

Related literature There is a growing empirical literature on the effect of acquisitions on innovation. [Cunningham *et al.* \(2021\)](#) show that in the US pharmaceutical industry, acquirers are more likely to stop drug research projects of acquired firms when these overlap with their own drug portfolio. These killer acquisitions are more frequent for acquirers with a dominant market position. [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

⁶This common negative effect explains the fall in the percentage of implemented startup ideas, despite the positive partial equilibrium effect uncovered by our regressions.

firms aiming to be acquired. Using data on publicly traded firms, they show that small firms' R&D increases after an acquisition shock. [Bena and Li \(2014\)](#), [Kim \(2020\)](#) and [Liu \(2022\)](#) study the effect of M&As on innovation and knowledge spillovers. 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](#); [Brutti and Rojas, 2021](#); [Eisfeld, 2023](#)). These studies are based on partial equilibrium models, while our contribution is to provide an aggregate general equilibrium perspective.

In the macroeconomic literature, [Jovanovic and Rousseau \(2002\)](#), [Dimopoulos and Sacchetto \(2017\)](#) and [David \(2020\)](#) analyze the effects of M&As on the allocation of capital, but do not consider innovation and productivity growth.⁷ More closely related to us, [Cavenaile, Celik and Tian \(2021\)](#) develop an endogenous growth model with mergers between incumbents. Our focus is different: we study startup acquisitions, leading us to consider novel issues such as the effects of acquisitions on the startup rate and on the implementation of ideas.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](#); [Akcigit and Kerr, 2018](#); [Peters, 2020](#)), by extending its standard framework to incorporate acquisitions and study their macroeconomic impact.

The remainder of the paper is organized as follows. Section 2 presents our model and highlights the channels through which startup acquisitions affect growth. Section 3 describes our micro data, lays out stylized facts, and empirically estimates the effects of acquisitions on startup ideas. Section 4 presents our calibration and our main quantitative results. Section 5 discusses extensions and robustness checks, and Section 6 concludes.

⁷[Pellegrino \(2022\)](#) and [Cao and Zhu \(2022\)](#) instead analyze the macroeconomic effect of M&As on markups. There is also an extensive literature on the microeconomic effects of M&As, including [Rhodes-Kropf and Robinson \(2008\)](#), [Andersson and Xiao \(2016\)](#), [Blonigen and Pierce \(2016\)](#) and [Wollmann \(2019\)](#).

⁸Likewise, [Weiss \(2022\)](#) uses an endogenous growth model to study changes in innovation costs, but does not consider the effect of acquisitions on the startup rate and on the implementation of ideas. [Pearce and Wu \(2022\)](#) study growth and welfare effects of acquisitions, but their focus is on brands of incumbent firms (with an application to the retail sector), and not startups. Our paper is also related to [Celik, Tian and Wang \(2022\)](#), who study the effects of information frictions in the merger market on innovation and business dynamism.

2 Model

In this section, we develop a model of the linkages between startup acquisitions and innovation. While we build on Schumpeterian heterogeneous-firm growth models (particularly Peters, 2020), our model introduces two new elements: a distinction between the invention and the implementation of ideas, and the possibility of startup acquisitions.

2.1 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

$$U = \int_0^{+\infty} \exp(-\rho \cdot t) \cdot \ln(C_t) dt, \quad (1)$$

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 $L > 0$ units of time, which she supplies inelastically at the wage w_t . The household owns all firms in the economy and accumulates wealth \mathcal{A}_t according to the budget constraint $\dot{\mathcal{A}}_t = r_t \cdot \mathcal{A}_t + w_t \cdot L - C_t$, where r_t is the rate of return on assets.

The final good is produced under perfect competition and assembled from a continuum of differentiated products with a Cobb-Douglas production function. Precisely,

$$Y_t = \exp\left(\int_0^1 \omega_{jt} \cdot \ln\left(\frac{y_{jt}}{\omega_{jt}}\right) dj\right), \quad (2)$$

where y_{jt} is the output of product j at instant t and ω_{jt} is the quality of product j at instant t . Product quality can take values in a finite set Ω , and products transition from state ω to state ω' at an exogenous Poisson rate $\tau_{\omega, \omega'}$. We assume that the economy starts in the steady state of this process and normalize $\int_0^1 \omega_{jt} dj = 1$.

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

$$y_{jft} = a_{jft} \cdot l_{jft}, \quad (3)$$

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 static Bertrand competition on product markets, implying that each product is only produced by the highest-productivity firm in equilibrium. We denote the productivity of this firm by a_{jt} . 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 implementation in order to develop these ideas. The next section describes these technologies.

Research and Implementation Innovations are generated by incumbent firms (which already produce at instant t) and by a large mass of startups (potential entrants).

To generate an idea at a Poisson arrival rate z , an incumbent must pay a research cost of $\zeta_I \cdot z^\psi \cdot Y_t$ units of the final good. In this cost function, $\zeta_I > 0$ is a scaling factor 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 aggregate output, to ensure balanced growth.

To develop an idea, the incumbent needs to invest into implementation. Precisely, if the incumbent invests $\kappa_I \cdot i_I^\psi \cdot Y_t$ units of the final good (with $\kappa_I > 0$), it successfully implements the idea with probability i_I .⁹ 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 immediately implemented disappears forever.

Ideas and innovations are also generated by startups. A startup can be created at a fixed cost $\zeta_S \cdot Y_t$ and generates a Poisson arrival rate 1 of ideas. A startup's idea applies to a randomly drawn good $j \in [0, 1]$. When the startup invests $\kappa_S \cdot i_S^\psi \cdot Y_t$ units of the final good (with $\kappa_S > 0$), it implements the idea with probability i_S . As with incumbent ideas, startup ideas are either implemented immediately or never. To allow for the empirical fact that startup ideas on average represent larger advances than incumbent ideas (as we will show in Section 3.2), we assume that a startup idea increases productivity by $n_S = 1 + N$ steps, 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. The quality of the idea is only revealed after investing into implementation.

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 invest into implementation: 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.

⁹To be exact, we assume that the implementation probability is $\min(i_I, 1)$. However, we choose parameter values ensuring that probabilities are always well below 1. For simplicity, we therefore omit the min operator. The same statement applies to all other implementation and meeting probabilities introduced below.

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 $\chi \cdot s^\varphi \cdot Y_t$ (with $\chi > 0$ and $\varphi > 1$) units of the final good in order to generate a probability s to meet any startup that has an idea on the incumbent's product. Thus, search costs are increasing and convex in the search effort. As usual, they also scale with aggregate output 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 order to monitor the market.

When there is a meeting, the incumbent may acquire the startup. Then, the incumbent transfers p_{jt}^A 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 acquisition price is determined through Nash bargaining, where the incumbent has a bargaining weight $\alpha \in (0, 1)$. The incumbent then invests to implement the startup's idea, using its own implementation technology. That is, by investing $\kappa_I \cdot i_A^\psi \cdot Y_t$ units of the final good, it implements the startup's 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. 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 (if the latter has lower implementation costs, i.e., if $\kappa_I < \kappa_S$). Second, acquisitions prevent entry, and therefore prevent the destruction of incumbent rents. While the first source of surplus represents a socially valuable transfer of ideas, the second does not.

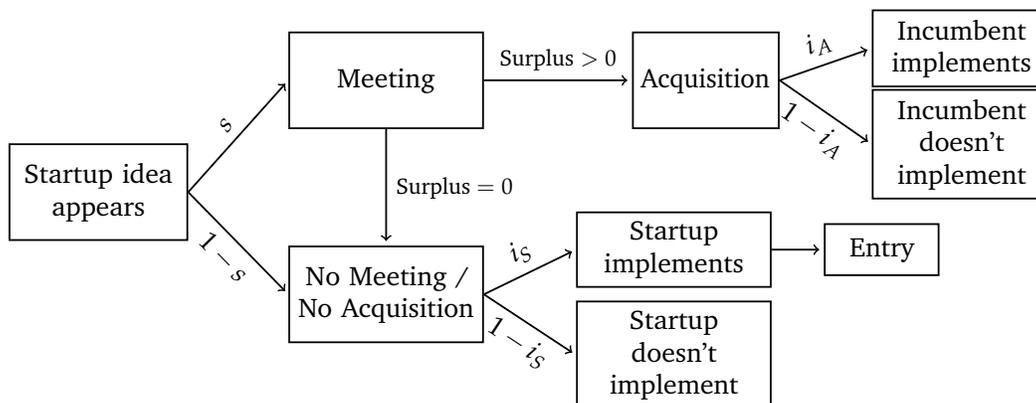


Figure 1: Timing of events for a startup idea within a period $(t, t + dt)$.

Life-cycle of a startup idea Figure 1 summarizes the timing of events for a startup idea within an instant of length dt . After the idea appears, the incumbent notices it and meets the

startup with probability s . If there is no meeting, the startup invests into implementation, which leads to possible entry and displacement of the incumbent (with probability i_S). If there is a meeting and the acquisition surplus is positive, an acquisition occurs and the incumbent chooses the probability i_A with which to implement the startup's idea.

2.2 Equilibrium

2.2.1 Household decisions, prices and profits

Throughout, we consider a balanced growth path (BGP) equilibrium with positive startup creation, in which all aggregate variables grow at a constant rate g . On the BGP, consumption growth holds the Euler equation

$$\frac{\dot{C}_t}{C_t} \equiv g = r - \rho. \quad (4)$$

Bertrand competition implies that each product is only produced by the highest-productivity firm. However, pricing decisions depend on the relative productivity of this incumbent with respect to the firm with the second-highest productivity. In our model, this “follower” is an old incumbent, displaced by the current incumbent when the latter overtook production. Denoting by a_{jt}^F the productivity of the follower firm for product j at instant t , we define the “technology gap” (the number of productivity steps between the incumbent and the follower) as the integer n_{jt} holding

$$\lambda^{n_{jt}} \equiv \frac{a_{jt}}{a_{jt}^F}, \quad (5)$$

The demand for each product j holds $y_{jt} = \omega_{jt} \cdot Y_t / p_{jt}$. As the price elasticity equals 1, an unconstrained monopolist would choose an arbitrarily high price. However, the price of any incumbent 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, prices hold

$$p_{jt} = \mu(n_{jt}) \cdot \frac{w_t}{a_{jt}} \quad \text{with} \quad \mu(n_{jt}) = \lambda^{n_{jt}}, \quad (6)$$

and profits are given by

$$\pi_t(\omega_{jt}, n_{jt}) = \omega_{jt} \cdot (1 - \lambda^{-n_{jt}}) \cdot Y_t. \quad (7)$$

Equation (7) shows that profits are increasing in product quality ω_{jt} , and increasing and concave in the technology gap n_{jt} . They do not depend on the productivity level a_{jt} .

2.2.2 Research, implementation 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 implementation probability, and whenever they meet a startup, they must decide whether or not to acquire it.

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

$$\begin{aligned}
r \cdot V_t(\omega, n) = & \max_{z,s} \left\{ \underbrace{\pi_t(\omega, n)}_{\text{Profits}} - \underbrace{\zeta_I \cdot z^\psi \cdot Y_t}_{\text{Research cost}} - \underbrace{\chi \cdot s^\varphi \cdot Y_t}_{\text{Search effort}} \right. \\
& + z \cdot \underbrace{\max_{i_I} \left[i_I \cdot \left(V_t(\omega, n+1) - V_t(\omega, n) \right) - \kappa_I \cdot i_I^\psi \cdot Y_t \right]}_{\text{Own innovation}} \\
& \left. + x \cdot \underbrace{\left[s \cdot V_t^{\text{Meet}}(\omega, n) + (1-s) \cdot V_t^{\text{NoMeet}}(\omega, n) - V_t(\omega, n) \right]}_{\text{Startup appears}} \right\} \\
& + \underbrace{\sum_{\omega' \in \Omega} \tau_{\omega, \omega'} \cdot \left[V_t(\omega', n) - V_t(\omega, n) \right]}_{\text{Quality shock}} + \underbrace{\dot{V}_t(\omega, n)}_{\text{Drift}}. \tag{8}
\end{aligned}$$

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 an implementation probability i_I . An implemented idea increases its technology gap by one step. 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 $V_t^{\text{Meet}}(\omega, n)$ the expected continuation value of the incumbent if there is a meeting, and by $V_t^{\text{NoMeet}}(\omega, n)$ the expected continuation value if 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

$$V_t^{\text{NoMeet}}(\omega, n) = \left[1 - i_{S,t}(\omega, n) \right] \cdot V_t(\omega, n), \quad (9)$$

where $i_{S,t}(\omega, n)$ is the optimal implementation probability chosen by the startup. When the startup does not implement, the incumbent's continuation value is just its current value. Instead, when the startup implements, 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 $V_{S,t}^{\text{NoMeet}}(\omega)$. This quantity holds

$$V_{S,t}^{\text{NoMeet}}(\omega) = \max_{i_S} \left\{ i_S \cdot \left(\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot V_t(\omega, n_S) \right) - \kappa_S \cdot i_S^\psi \cdot Y_t \right\} \quad (10)$$

where $\theta(n_S) \equiv \exp(-\gamma) \cdot \frac{\gamma^{n_S-1}}{(n_S-1)!}$ is the Poisson probability that the startup's innovation advances productivity by n_S steps. In the absence of a meeting, a startup chooses an optimal implementation probability i_S . When its idea is implemented, the startup becomes the new incumbent. With probability $\theta(n_S)$, it takes n_S steps on the productivity ladder, yielding a technology gap of n_S over the previous incumbent (who is now the follower). Instead, if implementation fails, the startup 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 $\Sigma_t(\omega, n)$. The surplus holds

$$\Sigma_t(\omega, n) = \max \left[0, \max_{i_A} \left\{ V_t(\omega, n) + i_A \cdot \sum_{n_S=1}^{+\infty} \theta(n_S) \cdot \left(V_t(\omega, n + n_S) - V_t(\omega, n) \right) - \kappa_I \cdot i_A^\psi \cdot Y_t \right\} - V_t^{\text{NoMeet}}(\omega, n) - V_{S,t}^{\text{NoMeet}}(\omega, n) \right]. \quad (11)$$

In equation (11), 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 $V_t(\omega, n)$. Moreover, the incumbent acquires the startup's idea and chooses an optimal

implementation probability i_A for 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. Finally, the incumbent transfers the acquisition price to the startup. This acquisition price is the startup's post-acquisition value. Thus, it is a pure transfer between the two firms and does not feature in the joint value shown in equation (11). To obtain the surplus, we subtract from the joint value the outside options of incumbent and startup, i.e., their expected values in the absence of a meeting.

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

$$V_t^{\text{Meet}}(\omega, n) = V_t^{\text{NoMeet}}(\omega, n) + \alpha \cdot \Sigma_t(\omega, n). \quad (12)$$

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

$$V_{S,t}^{\text{Meet}}(\omega, n) = V_{S,t}^{\text{NoMeet}}(\omega, n) + (1 - \alpha) \cdot \Sigma_t(\omega, n). \quad (13)$$

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

$$\zeta_S \cdot Y_t = \mathbb{E}_t \left[s_t(\omega, n) \cdot V_{S,t}^{\text{Meet}}(\omega, n) + \left(1 - s_t(\omega, n)\right) \cdot V_{S,t}^{\text{NoMeet}}(\omega, n) \right], \quad (14)$$

where $s_t(\omega, n)$ denotes the search effort by an incumbent with quality ω and technology gap n . This equation shows that the cost of creating a startup, $\zeta_S \cdot Y_t$, must be equal to the expected benefit of creating a startup. To compute the latter, note that the startup's idea falls on a random product j , characterized by a quality ω and a technology gap n . The expectation operator on the right-hand side of equation (14) refers to the joint distribution of products over these states. Depending on whether the startup meets an incumbent or not, it obtains one of the continuation values defined in equations (10) and (13).

Optimal policies To solve for the BGP policies, we guess and verify that the incumbent's value function holds $V_t(\omega, n) = v(\omega, n) \cdot Y_t$, i.e., that value scales one-to-one with aggregate output. This allows us to solve for the model in two blocks: we can first (numerically) solve for the startup rate, innovation and acquisition decisions, and the invariant distribution of products across quality levels and technology gaps. Then, in a second step, we can compute wages, consumption and the growth rate of the economy.

Appendix A.1 rewrites the HJB equation using the normalized value function v , and derives analytic expressions for the optimal research investment rate $z(\omega, n)$ and the optimal search effort $s(\omega, n)$ by incumbents. In both cases, incumbents equalize the marginal cost and the marginal benefit of these investments. For research, the marginal benefit is the expected gain in incumbent firm value from discovering (and potentially implementing) an own idea. For search, the marginal benefit is the expected surplus that the incumbent can extract from an acquisition.

Table 1, in turn, summarizes the optimal implementation probabilities for startup ideas implemented by the startup (i_S) or the incumbent (i_A), and for incumbent's own ideas (i_I). Incumbents and startups may make different implementation choices for the same idea. These differences stem from differences in both marginal costs and marginal benefits. First, implementation costs may differ: all else equal, a lower marginal cost (a lower cost shifter κ) implies higher investment. Second, 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 (as captured by the term $v(\omega, n)$ subtracted from the marginal benefit in the second and third lines of Table 1). When this last effect dominates, some ideas that would have been implemented by a startup will not be implemented by the incumbent, i.e., some acquisitions will be killer acquisitions.

Table 1: Optimality conditions for implementation probabilities

Implementer	Idea Creator	Marginal cost	Marginal benefit
Startup	Startup	$\kappa_S \cdot \psi \cdot (i_S)^{\psi-1}$	$= \sum_{n_S=1}^{+\infty} \theta(n_S) \cdot v(\omega, n_S)$
Incumbent	Startup	$\kappa_I \cdot \psi \cdot (i_A)^{\psi-1}$	$= \left(\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot v(\omega, n + n_S) \right) - v(\omega, n)$
Incumbent	Incumbent	$\kappa_I \cdot \psi \cdot (i_I)^{\psi-1}$	$= v(\omega, n + 1) - v(\omega, n)$

As we show in Appendix A.2, the optimal research, implementation, search and acquisition choices, together with the startup rate x , pin down the invariant distribution of products across quality and technology gaps. Appendix A.3 shows how we solve (numerically) for all these outcomes, using a simple algorithm. Crucially, this solution is independent of wages, consumption, and of the growth rate. We determine these variables in a second step, as shown in the next section.

2.2.3 Aggregate outcomes

To close the model, we need to solve for wages, consumption and aggregate growth. Using the demand function, it is easy to show that the incumbent producing product j

demands $\frac{\omega}{\mu(n)} \cdot \frac{Y_t}{w_t}$ units of labor. Imposing labor market clearing, we obtain the aggregate labor share:

$$\frac{w_t L}{Y_t} = \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} \left(m(\omega, n) \cdot \frac{\omega}{\mu(n)} \right), \quad (15)$$

where $m(\omega, n)$ stands for the mass of incumbents with quality ω and technology gap n .

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

$$Y_t = C_t + R_t + I_t + S_t. \quad (16)$$

In Appendix A.4, we show that R_t , I_t and S_t are linear functions of aggregate output. Therefore, aggregate consumption C_t grows at the same rate as output. Still in Appendix A.4, we show that this common growth rate is given by

$$g = \ln(\lambda) \cdot \sum_{\omega \in \Omega} \left(\sum_{n=1}^{+\infty} m(\omega, n) \cdot \omega \cdot \left[b_I(\omega, n) + b_S(\omega, n) \cdot (1 + \gamma) \right] \right), \quad (17)$$

where b_I is the arrival rate of innovations from incumbent ideas, and b_S is the arrival rate of innovations from startup ideas. These quantities are given by

$$\begin{aligned} b_I(\omega, n) &\equiv z(\omega, n) \cdot i_I(\omega, n) \\ b_S(\omega, n) &\equiv x \cdot \left(s(\omega, n) \cdot i_A(\omega, n) + (1 - s(\omega, n)) \cdot i_S(\omega) \right). \end{aligned} \quad (18)$$

The growth rate depends on the aggregate innovation rate. This rate is an average of innovations by incumbents and startups, weighted by product quality ω (which in equilibrium equals a product's sales share) and by the average step size of ideas.

2.3 Key properties of the model

Having characterized the BGP equilibrium, we now discuss some of its key properties.¹⁰ Figure 2 plots the normalized value function v and the research policy function z for incumbents. 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. The research investment of the firm, in turn, depends on the increments of the value function. Therefore,

¹⁰All figures in this section are produced using our baseline calibration, discussed in Section 4.1.

it is increasing in quality ω and decreasing in the technology gap n . Note that a firm continues to invest into research no matter how high its markup becomes, though research incentives become arbitrarily small for firms with very high markups.

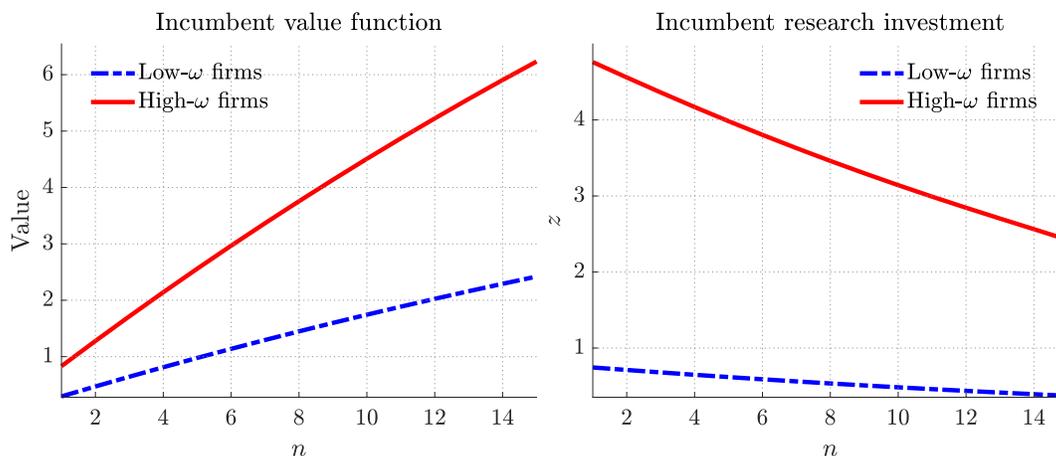


Figure 2: Value functions and research policy functions of incumbent firms, by firm type.

Figure 3 plots the (normalized) acquisition surplus and the incumbent meeting probabilities s . Acquisitions generate a surplus for two reasons. First, when incumbents have lower implementation costs ($\kappa_I < \kappa_S$), they transfer an idea to a more efficient user. 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. Some of these rents are kept by the incumbent, and the remainder is transferred to the startup. A higher quality ω and a higher technology gap n imply greater benefits of idea transfers and greater rents. Thus, the acquisition surplus is increasing in both variables, and firms with higher quality and higher technology gaps invest more resources into searching for startups.

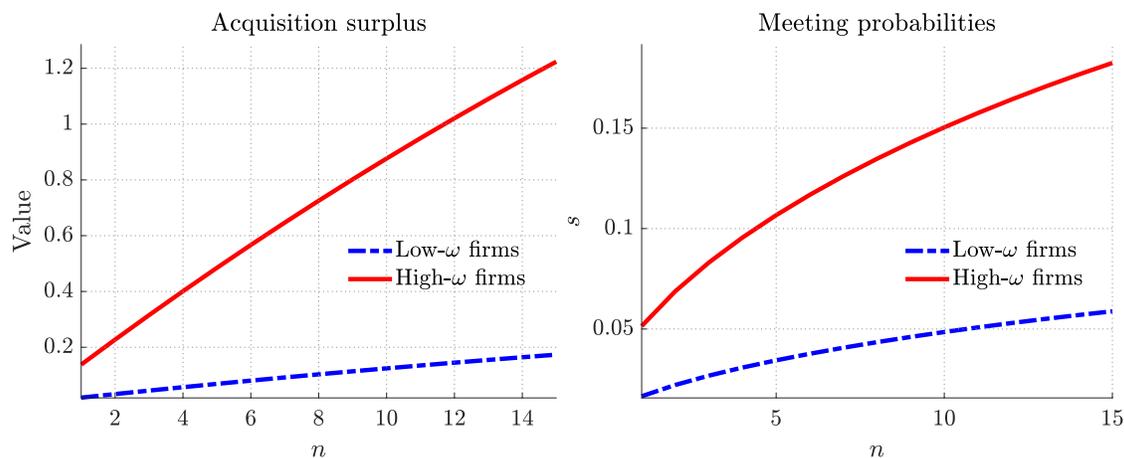


Figure 3: Acquisition surplus and meeting probabilities, by firm type.

Finally, the left panel of Figure 4 plots the implementation probabilities for a startup idea, distinguishing between the case in which the startup is not acquired and invests into implementation itself, and the case in which the startup is acquired and the incumbent invests into implementation. In this figure (and in our baseline calibration), incumbents have 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 a consequence, incumbents start to choose lower implementation probabilities than startups for the same idea. Therefore, acquisitions by dominant firms will on average “kill” some startup ideas.

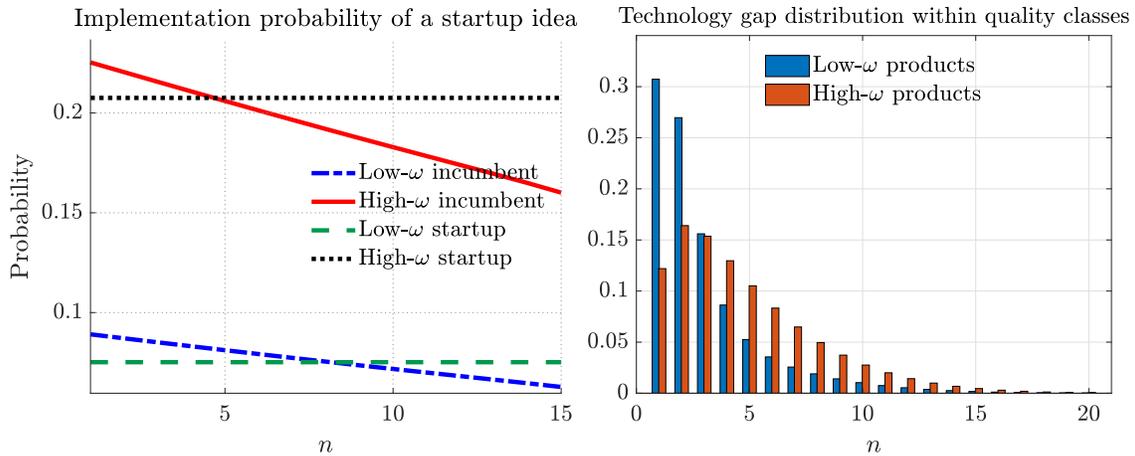


Figure 4: Implementation probabilities by firm type, and the invariant distribution of technology gaps.

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 products, shown in the right panel of Figure 4, is an important equilibrium object. This distribution is endogenous, shaped by the innovation choices of incumbents and startups. In particular, as the figure shows, high-quality products have on average higher technology gaps, as they receive more innovation.

Changes in public policy regarding startup acquisitions will affect all endogenous variables in our model (including the distribution). To organize our discussion of these effects, the next section shows a useful decomposition result.

2.4 Decomposing the effect of startup acquisitions on growth

Our paper aims to quantify the effects of startup acquisitions on growth. To do so, we will compare BGP equilibria of our model obtained for different values of search costs or

different antitrust policies. Using our expression for the growth rate in equation (17), we can show that the change in the growth rate between an baseline BGP equilibrium A and an alternative BGP equilibrium B can be expressed as:

$$\frac{g^B}{g^A} = \text{Share}_{\text{inc}}^A \cdot \frac{\text{Incumbent own innovation}^B}{\text{Incumbent own innovation}^A} + (1 - \text{Share}_{\text{inc}}^A) \cdot \left(\frac{\text{Startup rate}^B}{\text{Startup rate}^A} \cdot \frac{\text{Perc. of impl. startup ideas}^B}{\text{Perc. of impl. startup ideas}^A} \right), \quad (19)$$

where $\text{Share}_{\text{inc}}^A$ stands for the share of growth accounted for by incumbents' own innovation in the baseline BGP.¹¹

Equation (19) shows that in order to quantify the effect of any change in public policy regarding startup acquisitions, it is sufficient to know the response of the three margins shown in the equation: (i) changes in incumbent's own innovation (a sales-weighted average of the innovation rates of all incumbents), (ii) changes in the startup rate, and (iii) changes in the sales-weighted percentage of startup ideas that are implemented. Changes in these margins then need to be weighted by the baseline BGP share of growth accounted for by incumbents' own innovation.

All three margins shown in equation (19) are affected by shocks to startup acquisitions. Consider, for instance, a ban on startup acquisitions, a policy that we will formally analyze in Section 4. Such a ban clearly affects the incentives for startup creation, the likelihood that startup ideas get implemented (especially if startups and incumbents implement at different rates), and the incentives for incumbents' own innovation. The aim of our paper is to quantify these effects, to get at the overall effect of startup acquisitions on growth.

In order to achieve this objective, we need to map the model to the data. Thus, in the next section, we analyze micro data on innovation and startup acquisitions. This data shows how frequent startup acquisitions are, which incumbents are more likely to acquire startups, and allows us to explore the effects of these operations on the involved firms and ideas. We will then use this information to calibrate our model's parameters.

¹¹Equation (19) is formally derived in Appendix A.5. The variables in the decomposition are given by

$$\begin{aligned} \text{Startup rate} &= x \\ \text{Incumbent own innovation} &= \sum_{\omega, n} m(\omega, n) \cdot \omega \cdot b_I(\omega, n) \\ \text{Perc. of impl. startup ideas} &= \sum_{\omega, n} m(\omega, n) \cdot \omega \cdot \left(s(\omega, n) \cdot i_A(\omega, n) + (1 - s(\omega, n)) \cdot i_S(\omega) \right) \\ \text{Share}_{\text{inc}}^A &= \ln(\lambda) \cdot \text{Incumbent own innovation}^A / g^A \end{aligned}$$

3 Stylized facts and empirical analysis

3.1 Startup acquisitions in the United States

To begin with, we need to get a sense of the overall importance of startup acquisitions in the United States. This is not straightforward, as there is a limited amount of publicly available data on startup activity. The most comprehensive data collection effort is due to [Guzman and Stern \(2020\)](#), who compiled a database containing all new firms incorporated in 32 states.¹² Their data contains information about firm characteristics at incorporation (e.g., whether the firm holds a patent application) and about the firm’s growth outcomes. In particular, for all firms incorporated between 1988 and 2008, [Guzman and Stern \(2020\)](#) record whether these firms are acquired, do an initial public offering (IPO), or grow to 100 or more employees during their first six years of existence.

Table 2: Growth outcomes for newly incorporated firms.

Sample	(1) All new firms	(2) Patenting new firms
Total number	18,764,856	37,588
<i>Outcome after 6 years</i>		
... Acquisition	0.06%	4.02%
... IPO	0.01%	1.13%
... 100+ employees	0.23%	6.60%

Source: [Guzman and Stern \(2020\)](#) and own computations. The sample contains all newly incorporated firms incorporated in 32 US states between 1988 and 2008. Column (1) refers to all new firms, and column (2) to new firms with a patent application.

Column (1) in Table 2 shows that in the overall population of new firms, acquisitions are very rare events: only 0.06% of firms are acquired within their first 6 years of existence. However, it is well known that most newly created firms do not have growth ambitions and remain small throughout their existence ([Hurst and Pugsley, 2011](#)). Thus, for our purpose, it is more relevant to consider a subsample of potentially innovative and growth-oriented new firms, comparable to the startups in our model. To do so, we focus on firms that hold a patent application at the time of incorporation. As shown in column (2), acquisitions of these patenting startups are much more frequent: about 4% of them are acquired within their first six years of existence. This number will be an important calibration target for

¹²These states represent 80% of US GDP. Data can be downloaded at <https://www.startupcartography.com/>.

our model. Patenting startups are also much more likely to achieve an IPO or significant employment growth.¹³

Figure 5 plots the percentage of acquired patenting startups by incorporation year. Acquisitions peak for the 1999-2000 startup cohorts, in the middle of the dot-com boom, at around 6%. However, there does not appear to be a decisive trend over time.¹⁴ This is in line with our focus on a BGP equilibrium, in which this percentage is constant.

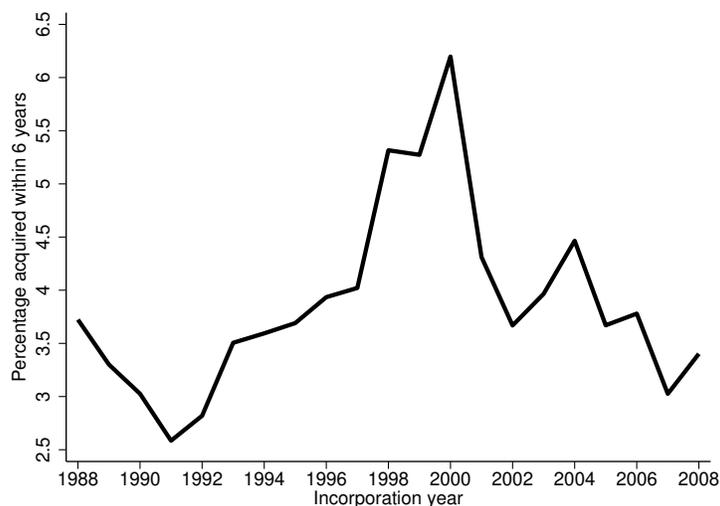


Figure 5: Acquisitions of patenting startups (data from [Guzman and Stern, 2020](#)).

These facts provide us a good first sense of the frequency of startup acquisitions. However, the [Guzman and Stern \(2020\)](#) database does not contain information about the acquiring firm or about the startup's patenting behavior (beyond the fact of holding a patent application in the incorporation year). Therefore, in the next section, we construct a data set that includes these elements. While our data has some disadvantages with respect to [Guzman and Stern](#) (most importantly, the fact that we do not observe firm incorporation dates), it also contains important new information, allowing us to describe further characteristics of startup acquisitions and, most importantly, to provide some causal evidence of the effect of startup acquisitions on the involved firms.

¹³Table B.5 in the Appendix shows that acquisitions, IPOs and employment growth are even more prevalent among patenting firms that, apart from incorporating in their home state, also file an incorporation in Delaware (which offers tax and judicial advantages). Indeed, [Guzman and Stern \(2020\)](#) show that holding a patent and incorporating in Delaware is the strongest predictor of entrepreneurial success.

¹⁴In line with the literature, we do observe a downward trend in the percentage of startups doing an IPO (see e.g. [Ewens and Farre-Mensa, 2020](#)), as well as in the percentage of startups experiencing strong employment growth (see e.g. [Decker et al., 2016](#)). The fact that startup acquisitions do not increase seems to indicate that these trends are not primarily due to high-growth startups being acquired more frequently.

3.2 Data, definitions and descriptives

3.2.1 Data construction

To construct our dataset, we merge three sources of information: data on acquisitions from the financial information provider Refinitiv (formerly Thomson Financial), patent data from the NBER Patent Data Project, and accounting data for public firms from Compustat. This section describes our data sources. Appendix B.1 contains further details.

Acquisitions To track acquisitions, we rely on the ThomsonONE database, using information between 1981 and 2014.¹⁵ The database provides transaction-level data on mergers and acquisitions (M&As) and includes practically all deals involving US firms over the considered time period. ThomsonONE provides several variables of interest, such as the names of the involved firms, the industries in which they operate, the announced and effective dates of the deal, the transaction value, and sometimes even the revenue and total assets of the involved firms.

Patents In order to measure the innovation activity of firms, we rely on patent data from the NBER Patent Data Project (NBER-PDP), which provides US patent data for 1976-2006.¹⁶ In addition to the patent owner, this data set also provides us with the forward and backward citations to the patent, a measure of each patent's originality and generality, and patent 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. To address this problem, the NBER-PDP standardizes commonly used words in firm names and provides a match to publicly listed firms (Bessen, 2009). However, further work is needed to match patents to private firms, as we explain below.

Company accounts Finally, we use the Compustat North America database, provided by Standard & Poor's.¹⁷ This database contains balance sheet and income statement information for all publicly traded firms in the United States.

Merging these three databases is straightforward for publicly listed firms, as both ThomsonONE and the NBER-PDP provide firm identifiers that are consistent with Compustat. For private firms, the situation is more challenging. For these firms, we do not have

¹⁵This is a commercial database, which can be accessed at <https://www.refinitiv.com/en/products/sdc-platinum-financial-securities>. Due to various changes for the providing firm, the database has frequently changed names and is currently branded as the Refinitiv SDC Platinum database. It is the standard database used in M&A analysis (see e.g. David (2020) or Guzman and Stern (2020)).

¹⁶The data set can be downloaded at <https://sites.google.com/site/patentdataprotect/>.

¹⁷<https://www.spglobal.com/marketintelligence/en/?product=compustat-research-insight>.

accounting data. Thus, in order to match them to their patents, we can only rely on their names. To carry out the match, we first standardize the company name provided by ThomsonONE. Then, we employ a fuzzy name matching algorithm and a large scale manual check to match each company to its patents, as described in greater detail in Appendix B.1.

3.2.2 Definitions and descriptive statistics

Startups: definition and importance Given our focus on innovative firms, we consider throughout acquisitions in which the acquired firm holds at least one patent. Our dataset does not allow us to observe the exact incorporation date of a firm, but it does provide us with its complete patenting history. Therefore, we define a firm as a startup if it is within 6 years of its first patent.

Our data is consistent with startups being important for aggregate innovation. Startups account for 27% of all patent applications. Remarkably, however, their patents collect 74% of all patent citations. This suggests that startup patents are on average of higher quality than patents filed by older firms (in line e.g. with [Akcigit and Kerr, 2018](#)). Thus, startup innovation (and the way in which it is affected by acquisitions) is likely to have a large impact on aggregate outcomes.

Regarding acquisitions, we find that around 49% of all transactions with a publicly listed acquirer and a private target in our sample are acquisitions of patenting startups. Thus, innovative startup targets account for a sizeable share of overall acquisition activity. Table B.1 in the Appendix shows some summary statistics for our acquisition sample.

Selection into acquisition The acquisition process is obviously not random: both the acquiring firms and the startups that they acquire are a selected sample of the overall population of firms. For instance, Figure 6 compares, for any given year, the average sales of publicly listed firms that acquire startups to the average sales of publicly listed firms that do not acquire startups. Throughout our sample period, acquirers are systematically larger than non-acquirers, by a factor of about 2.6. This implies that the average acquirer is by a factor 2.1 larger than the average firm. Our model also suggests this link (as larger firms have higher rents to protect), and we will target this number in our calibration.

Likewise, acquired startups are different from non-acquired startups. For instance, we find that the patents of an acquired startup are cited four times as much as the average startup patent. In sum, there appears to be positive assortative matching in the acquisition process, as the largest incumbents match with the “best” startups.

These stylized facts provide some further information on the importance and characteristics of startup acquisitions. However, our data also allows us to dig deeper into the effects

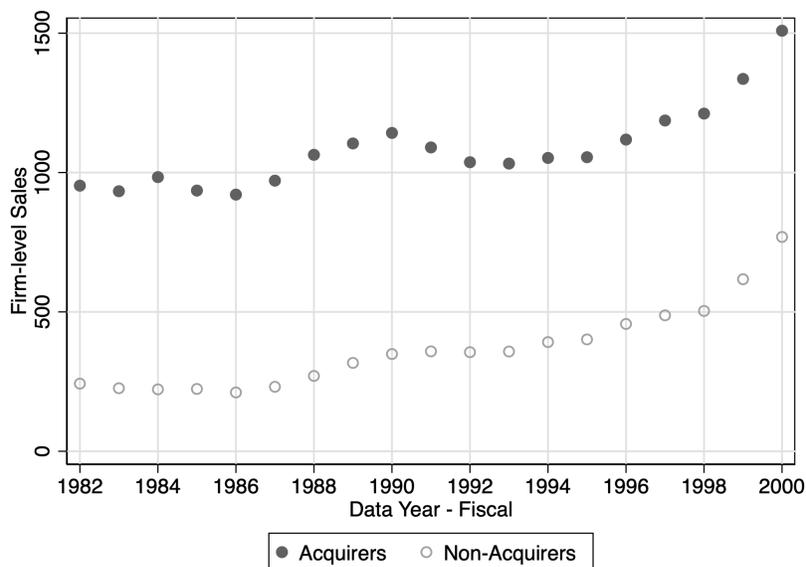


Figure 6: Sales by type of firm: acquirers vs. non-acquirers. Data from ThomsonONE, NBER Patent Data project and Compustat.

of these acquisitions on innovation. We do so in the next section.

3.3 The effect of acquisitions on the implementation of ideas

As we have shown in Section 2, one important channel through which acquisitions affect growth is their effect on the implementation probability of startup ideas. An acquisition may increase this probability (if incumbents have advantages in developing ideas and bringing them to the market) or decrease it (if incumbents are engaging in killer acquisitions). In this section, we try to assess the relative strength of these forces.

To do so, we need a proxy for the implementation of startup ideas. We propose to rely on the evolution of patent citations after the acquisition event. That is, we consider the set of patents that the startup held before the acquisition. If citations to these pre-existing patents increase after the acquisition, we interpret this as evidence for the startup's ideas being further developed and built upon. If, on the other hand, citations to these patents decrease after the acquisition, we interpret this as evidence for the idea being shelved.¹⁸

Of course, just considering the change in patent citations after acquisition faces an endogeneity problem: in the previous section, we have shown that acquired patents are different from the average patent. Therefore, we use a nearest neighbor matching algorithm

¹⁸In line with this interpretation, [Argente, Baslandze, Hanley and Moreira \(2020\)](#) show that in the consumer goods sector, more highly cited patents lead to a higher likelihood of introducing new products.

to link each treated patent (belonging to a startup that will eventually be acquired) to up to ten control patents (belonging to non-acquired startups). We match on several patent and firm characteristics, including patent application year, technological subsector and a number of measures of patent novelty, originality and impact. We artificially assign to each control patent the acquisition year of its matched treated patent. Appendix B.2 discusses the matching method in greater detail. In particular, Table B.2 shows that there are no significant differences between treated and control patents at the time of the acquisition.

With this data, we run a difference-in-difference regression:

$$\begin{aligned} \text{NumCites}_{ijt} = & \beta_1 \cdot D(\text{Treatment})_i + \beta_2 \cdot D(\text{Post})_{it} \\ & + \beta_3 \cdot D(\text{Treatment})_i \cdot D(\text{Post})_{it} + \alpha_j + \alpha_t + u_{it}, \end{aligned} \quad (20)$$

where NumCites_{ijt} is the number of citations received by patent i belonging to patent pair j at year t , $D(\text{Treatment})_{ij}$ takes value 1 for treated patents, and $D(\text{Post})_{ijt}$ takes value 1 for the years after the acquisition. In our baseline analysis, we consider a 14-year window around the acquisition (i.e., 7 years before and 7 years after this event). Finally, α_j are matched patent pair fixed effects and α_t are year fixed effects.

Table 3: The effect of acquisitions on the implementation of ideas

Dependent variable: Number of citations received				
	(1)	(2)	(3)	(4)
D(Post)	0.405*** (0.028)	0.397*** (0.019)	0.439*** (0.030)	0.346*** (0.019)
D(Treatment)	-0.016 (0.068)	-0.014 (0.062)	-0.013 (0.038)	-0.010 (0.035)
D(Post) * D(Treatment)	0.228*** (0.051)	0.226*** (0.050)	0.222*** (0.044)	0.218*** (0.041)
Observations	206,432	206,432	206,352	206,352
Matched Pair FE			✓	✓
Year FE		✓		✓

Notes: We use a Poisson estimator. The dependent variable is the number of citations received at the patent-year level. $D(\text{Treatment})_i$ takes value 1 for treated patents, and $D(\text{Post})_{it}$ takes value 1 for the years after the acquisition. Standard errors are clustered at the target firm level. * significant at 10%; ** significant at 5%; *** significant at 1%.

In equation (20), a positive coefficient β_3 would imply that after being acquired, a treated patent receives relatively more citations (our proxy for the implementation of ideas) than a control patent. Instead, a negative coefficient β_3 would imply that a treated patent

receives relatively less citations after being acquired.

Table 3 presents the estimation results. In column (1) we show the results of a simple Poisson estimator with no fixed effects. The interaction term is positive and statistically significant, indicating that after an acquisition, the citations of treated patents increase by 22.8% relative to the citations of control patents. This magnitude is remarkably stable even after including year fixed effects (column 2), matched patent pair fixed effects (column 3), or both sets of fixed effects (column 4).¹⁹ Moreover, note that the estimate for the coefficient of the treatment dummy is always statistically indistinguishable from zero, indicating that treatment and control patents are equally cited before an acquisition event. Appendix B.4 discusses several additional robustness checks, including changes in the number of control patents, different event study windows, and sector-level regressions.

3.4 Heterogeneous effects

Our model predicts that the effect of an acquisition on a startup idea depends on the characteristics of the acquirer. For instance, an acquisition by a firm with a higher technology gap and thus a higher markup should have a less positive effect on implementation. Moreover, in Section 5, we will briefly consider an extended model, in which acquirers can buy both startups with an idea on their own product (implying a threat of displacement, as in the baseline model) and startups with ideas on other products. In this extended model, acquisitions of non-competing startups always have a more positive impact on implementation than acquisitions of competing startups.

We test these predictions in the data. Table 4 summarizes our results. In columns (1) - (2), we evaluate whether the effect of acquisitions depends on the acquirer's market share (which is likely to be positively correlated both with the technology gap and the markup, two objects that are hard to measure in the data). We split the sample by the median of the acquirer's sales share in its (SIC 3-digit) industry. Consistent with our prior, the magnitude of the estimated interaction term diminishes by almost half when the acquirer has a market share above the median (column 1) compared to the case when it has a market share below the median (column 2).

In columns (3)-(4) we empirically test the second hypothesis, i.e., whether competing directly with the target diminishes the acquirer's incentives to implement the target's idea. We generate a dummy taking value 1 if both acquirer and target firms have the same primary SIC 3-digit code. In our sample split exercise, we find that being active in the same

¹⁹A matched pair is defined by the bundle of one treated patent and up to ten control patents obtained from the matching algorithm based on a set of observables.

product market reduces the probability of implementing ideas developed by the target firm by more than half (column 3) compared to the case when the acquirer and the target firm are active in different product markets (column 4).

Table 4: Heterogeneous effects by market share and product market overlap

Dependent variable: Number of citations received

	Market Share		Same SIC3		Same SIC3/NAICS4	
	(1) Above	(2) Below	(3) Same	(4) Different	(5) Same	(6) Different
D(Post)	0.371*** (0.022)	0.334*** (0.020)	0.361*** (0.023)	0.332*** (0.021)	0.399*** (0.022)	0.325*** (0.022)
D(Treatment)	0.004 (0.052)	0.004 (0.044)	0.030 (0.045)	-0.043 (0.049)	0.004 (0.050)	-0.031 (0.046)
D(Post) * D(Treatment)	0.158*** (0.059)	0.249*** (0.048)	0.132** (0.054)	0.288*** (0.051)	0.158*** (0.055)	0.264*** (0.051)
Observations	88,187	92,480	83,500	122,817	67,359	130,598
Matched Pair FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓

Notes: We use a Poisson estimator. The dependent variable is the number of citations received at the patent-year level. $D(\text{Treatment})_i$ takes value 1 for treated patents, and $D(\text{Post})_{it}$ takes value 1 for the years after acquisition. Columns (1)-(2) split the sample by the median of acquirer market share defined at the SIC3-year level, where column (1) keeps the observations above the median in market share and column (2) the ones below. Columns (3)-(4) split the sample based on whether both acquirer and target have the same primary SIC 3-digit industry code. Finally, columns (5)-(6) replicate the exercise with a sample split requiring both firms to have the same SIC 3-digit industry code (until 1997) and the same NAICS 4-digit industry code (since 1997). Standard errors are clustered at the target firm level. * significant at 10%; ** significant at 5%; *** significant at 1%.

Finally, as a robustness exercise, we adjust the analysis of columns (3)-(4) to the fact that in 1997, the SIC industry classification was replaced by the NAICS industry classification. Therefore, our dummy now takes value 1 if acquirer and target are in the same SIC 3-digit industry until 1997 and in the same NAICS 4-digit industry after 1997. Results remain quantitatively similar (see columns (5)-(6)). Overall, results in this table show that the willingness to implement ideas developed by startups depends on the acquirer's characteristics, in line with our model.

Summing up, our findings in this empirical analysis indicate that the average acquisition increases the likelihood that a startup's idea is implemented. Through the lens of our model, this suggests that incumbents have lower implementation costs, and that this cost

advantage dominates the Arrow replacement effect. We will calibrate our model to match this fact, by matching the results of our main difference-in-difference regression.

However, the regressions presented in this section obviously miss general equilibrium effects that affect all firms equally. Moreover, as equation (19) shows, the implementation channel is not the only link between acquisitions and innovation: acquisitions also affect the innovation behavior of incumbents, as well as the incentives to create a startup. In the next section, we return to the model to jointly evaluate all these forces.

4 Quantitative analysis

4.1 Calibration strategy

We assume that a period of unit length in the model corresponds to one year in the data. We set some parameters externally. The discount rate is set to $\rho = 0.02$, which together with a 2% target for the annual growth rate implies a real interest rate of 4%. We assume that there are two product quality classes, $\Omega = \{\omega_L, \omega_H\}$ with $\omega_L < \omega_H$. At every point in time, 20% of firms belong to the H class, and their sales account for 80% of output (in line with the average industry-level sales share of the largest 20% of firms in Compustat). This implies $\omega_H/\omega_L = 16$. Firms transition from ω_H to ω_L at a Poisson rate $\tau = 0.1$, matching the fact that in every year, 10% of Compustat firms belonging to the top 20% of sales in their industry drop out of that category in the subsequent year. We set the elasticity of R&D costs to innovation to $\psi = 2$, following empirical evidence summarized in [Akcigit and Kerr \(2018\)](#). The average step size advantage for startup ideas is $\gamma = 0.415$. To obtain this number, we rely on [Kogan, Papanikolaou, Seru and Stoffman \(2017\)](#), who estimate that the elasticity of a patent’s market value to its number of forward citations is 0.17. Our results from Section 3.2 indicate that the average startup patent is cited 7.7 times as much as the average incumbent patent.²⁰ This implies that a startup patent is on average $7.7^{0.17} - 1 \approx 41.5\%$ more valuable than an incumbent patent, and we thus assume that the former represents on average 41.5% more steps. Finally, following [David \(2020\)](#), we set the Nash bargaining parameter for incumbents to $\alpha = 0.5$. In Section 5, we conduct extensive robustness checks around this baseline.

Our choices leave seven parameters to be identified: the productivity step size, λ ; the research and implementation cost shifters for incumbents, ζ_I and κ_I ; the fixed cost of startup creation, ζ_S ; the implementation cost shifter for startups, κ_S ; and the scale and curvature

²⁰Startups account for 27% of patents and for 74% of patent citations. Therefore, citations per patent are, on average, $\frac{0.74/0.27}{(1-0.74)/(1-0.27)} \approx 7.7$ times higher for startup patents than for non-startup (i.e. incumbent) ones.

parameters in the incumbent’s search cost function, χ and φ . We calibrate these parameters internally, using an indirect inference approach. That is, we choose values for the seven parameters that minimize the distance between seven model-generated moments and their empirical counterparts.²¹ The success of this calibration strategy relies on choosing moments that are relevant for the economic mechanisms we want to highlight, as well as sufficiently sensitive to variation in parameters. As the model is non-linear, all moments are affected by all parameters, making identification challenging. Nevertheless, we provide economic intuitions for the identification power of different moments. To support these intuitions, Appendix C.1 shows the results of a rigorous global identification test, indicating that all parameters are very well identified by the chosen moments.

Table 5: Calibrated parameters.

A. Externally Calibrated Parameters

Parameter	Description	Value	Target/Source
ρ	Discount rate	0.02	4% annual real interest rate
ω_H/ω_L	Relative product quality	16	Top 20% sales share (Compustat)
τ_{HL}	Transition rate from H to L quality	0.10	Likelihood to drop from Top 20% (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	0.415	Kogan <i>et al.</i> (2017) and Section 3

B. Internally Calibrated Parameters

Parameter	Description	Value
λ	Innovation step size	1.0295
ζ_S	Startup creation cost	0.0380
κ_S	Implementation cost scale for startups	2.4474
ζ_I	Research cost scale for incumbents	0.0038
κ_I	Implementation cost scale for incumbents	1.3905
χ	Search cost scale for incumbents	3.2137
φ	Search cost curvature for incumbents	2.7254

Table 5 lists the calibrated parameter values, and Table 6 shows the seven targeted moments. First, we target a growth rate of 2%, the long-run average growth rate of GDP per person in the United States (Jones, 2016). This target identifies the innovation step λ .

Second, we target the likelihood of acquisitions for startups. In Section 3.1, we found that around 4% of innovative startups in the United States are acquired. This target

²¹Formally, the vector of parameters $\theta = (\lambda, \zeta_I, \kappa_I, \zeta_S, \kappa_S, \chi, \varphi)$ is chosen to minimize the following criterion distance function: $\sum_{m=1}^7 \frac{|\text{Moment}_m(\text{Model}, \theta) - \text{Moment}_m(\text{Data})|}{0.5|\text{Moment}_m(\text{Model}, \theta)| + 0.5|\text{Moment}_m(\text{Data})|}$.

identifies χ , the search cost of incumbents for startups.

Third, we target the likelihood that an incumbent is displaced by an entrant, i.e., the exit rate of incumbents. To do so, we follow [Garcia-Macia *et al.* \(2019\)](#) and use data from the US Census Bureau’s Business Dynamics Statistics to compute an exit rate of “large” firms (with 20 employees or more). For the period 1988-2014, we find an exit rate of 7.3%.²² This target identifies ζ_S , the cost of startup creation. We also target the contribution of entrants to overall productivity growth. As this moment is not directly observable in our data, we use the estimate of [Akcigit and Kerr \(2018\)](#), who find that entry accounts for 25.7% of productivity growth in the United States.²³ This target identifies the research cost of incumbents, ζ_I , which shifts the growth contribution of incumbents’ own innovation.

Fourth, we use our regression evidence from Section 3.3 to set the relative implementation costs of startups and incumbents, κ_I/κ_S . In Table 3 we found that, on average, an acquisition increases the citation count of a startup patent by 22%. To make this finding operational, we need an assumption about the relationship between patent citations (the observed outcome in our empirical analysis) and the implementation probability of ideas (the observed outcome in our model). For this, we again rely on [Kogan *et al.* \(2017\)](#), who find an elasticity of 0.17 of patent value to patent citations. In our model, the expected value of an idea is equal to the product of the implementation probability and the value of the implemented idea (and the latter is ex ante identical for all startup ideas). Therefore, we assume that an acquisition increases the implementation probability of an idea by $0.17 \cdot 0.22 = 0.0374$ log points, and target this number in the model.²⁴ As shown in Panel B of Table 5, our calibration implies that incumbents have about 43% lower implementation costs than startups. Indeed, if costs were equal, the replacement effect would imply that acquisitions lower implementation. However, we observe an increase in implementation in the data, which the model rationalizes through lower costs for incumbents.

²²[Garcia-Macia *et al.* \(2019\)](#) use the Longitudinal Business Database (the confidential micro data underlying the BDS), and define large firms as those having more employees than average. Between 1988 and 2014, the average firm in the BDS had 22 employees. As the BDS only provides data aggregated by size classes (e.g., for firms between 10 and 19 employees), we choose the closest available cutoff of 20. As in [Garcia-Macia *et al.* \(2019\)](#), we compute exit rates over five-year intervals. They find an exit rate of 6% for large firms between 1983 and 2013, close to our baseline. Similarly, [Akcigit and Kerr \(2018\)](#) compute an entry rate for innovative firms of 5.8% (in the BGP of our model, the entry rate of innovative firms equals the exit rate of incumbents).

²³To obtain this number, [Akcigit and Kerr \(2018\)](#) structurally estimate a creative destruction model on the universe of patenting firms. [Garcia-Macia *et al.* \(2019\)](#), who focus on all firms, find similar numbers, namely a 21.1% contribution of entry to productivity growth over our sample period.

²⁴In the model, our empirical regression would yield a coefficient equal to the difference between the average implementation probability (in logs) in the sample of acquired startups, $\sum_{\omega} \sum_n m^A(\omega, n) \cdot \ln(i_A(\omega, n))$, and the average implementation probability in a control group of non-acquired startups, $\sum_{\omega} \sum_n m^A(\omega, n) \cdot \ln(i_S(\omega))$. Note that both averages need to be computed using the distribution of acquired startups over states (ω, n) , denoted $m^A(\omega, n)$. This distribution holds $m^A(\omega, n) = \frac{m(\omega, n) \cdot s(\omega, n)}{\sum_{\omega'} \sum_{n'} m(\omega', n') \cdot s(\omega', n')}$.

Fifth, we target an average implementation probability of non-acquired startups of 10%. This target pins down the level of startup implementation costs κ_S . While this statistic is hard to measure in the data, [Guzman and Stern \(2020\)](#) indicate that 6.6% of all innovative startups that are not acquired either achieve an IPO or grow to more than 100 employees. As we have previously identified incumbents as firms with 20 employees or more, we choose a higher target of 10%. However, as we show in [Appendix C.2](#), this target turns out to be irrelevant for all our quantitative results. That is, none of our results would change if we would target an average implementation probability of e.g. 15, 20 or 25%.²⁵

Finally, we target selection into acquisition (on the acquirer side), matching the empirical fact that acquirers' sales are on average 2.1 times larger than those of the average firm (see [Section 3.2](#)). This target identifies the parameter φ , the curvature in the search cost function, which governs how steeply search costs increase for firms that search harder.

Table 6: Model fit.

Targeted moment	Model	Data	Data source	Identifies
Growth rate	2.0%	2.0%	Jones (2016)	λ
Exit rate	7.3%	7.3%	BDS	ξ_S
Growth contribution of entrants	25.7%	25.7%	Akcigit and Kerr (2018)	ξ_I
Avg. implementation prob., startups	10.0%	10.0%	See text	κ_S
Effect of acq. on implementation prob.	0.0374	0.0374	Section 3	κ_I/κ_S
Percentage of startups acquired	4.0%	4.0%	Section 3	χ
Relative size of acquiring firms	2.10	2.10	Section 3	φ

As [Table 6](#) shows, the model matches all moments exactly. This is due to the tight link between targeted moments and parameters, shown more formally in [Appendix C.1](#). In [Section 5](#) and [Appendix C.2](#), we conduct extensive robustness checks around this baseline calibration.

4.2 The aggregate effects of startup acquisitions

We can now use the calibrated model to assess the effect of startup acquisitions on growth. To do so, we solve for the BGP equilibrium for different values of the search cost χ , keeping all other parameters at their baseline values. Recall that χ represents frictions in the acquisition market: a low value of this parameter implies low frictions and

²⁵Roughly speaking, growth depends on the own innovation rate of startups, which is pinned down by our targets for the exit rate and the contribution of entry to growth. Targeting the implementation probability of startups decomposes this innovation rate into the arrival rate of ideas and the probability that these are implemented, but this decomposition is irrelevant for aggregate outcomes. [Appendix C.2](#) explains this issue.

frequent acquisitions, while a high level implies high frictions and infrequent acquisitions. Accordingly, as shown in the left panel of Figure 7, the equilibrium frequency of acquisitions is monotonically decreasing in χ .

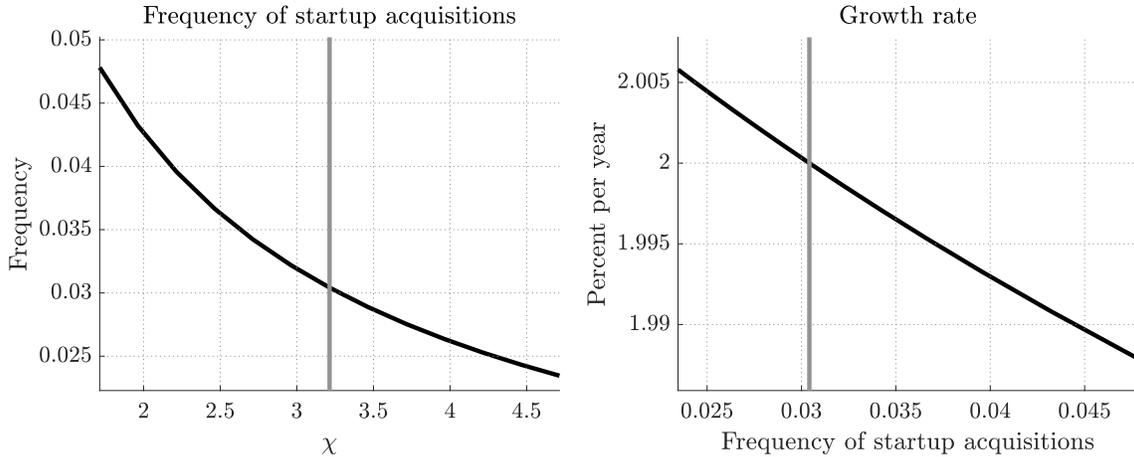


Figure 7: BGP equilibria for different values of search costs χ . *Notes:* The frequency of acquisitions is the product of the startup rate and the share of startups that are acquired. The baseline calibration value of χ is marked with a vertical line in the left plot. The right plot shows the reduced-form relationship between the frequency of acquisitions and growth. The vertical line marks the baseline frequency of acquisitions.

The right panel of Figure 7 plots the growth rate of the economy for different values of search costs. For convenience, we plot the growth rate directly against the frequency of acquisitions implied by different search costs.²⁶ This figure illustrates the main result of our paper: a higher frequency of startup acquisitions is associated with a lower growth rate.

Why is there a negative relationship between startup acquisitions and growth? To answer this question, we use our decomposition from equation (19), which showed that the change in the growth rate can be decomposed into changes in incumbent own innovation, changes in the startup rate and changes in the sales-weighted percentage of implemented startup ideas. The top left panel of Figure 8 plots changes in these three margins.

The figure shows, first of all, that more frequent acquisitions are associated with a higher startup rate. Indeed, in our model, acquisitions have a strong incentive effect on startup creation. As startups only sell to incumbents if the acquisition price exceeds their outside option of independent development, acquisitions always increase a startup’s payoff. Thus, all else equal, more frequent acquisitions must increase the value of creating a startup, which in equilibrium leads to more startup creation.

On its own, the incentive effect would suggest that startup acquisitions increase growth. However, it turns out that it is more than compensated by a decrease in incumbent’s own

²⁶This choice is made to improve readability. The frequency of acquisitions is obviously an endogenous outcome, and all variation in it is due to underlying variation in the search cost parameter χ .

innovation and in the percentage of startup ideas being implemented. The main reason for this decrease is the fact that the higher startup rate reduces the value of incumbents. Even though incumbents can acquire startups and thereby avoid displacement, an acquisition implies a costly sharing of rents with the threatening startup. Thus, a higher startup rate implies a higher frequency of displacement or rent sharing, two undesirable outcomes for incumbents. It is interesting to note that incumbents face a prisoner's dilemma of sorts: they would all gain by collectively agreeing not to acquire startups, but each of them has an incentive to deviate once it is threatened by displacement.

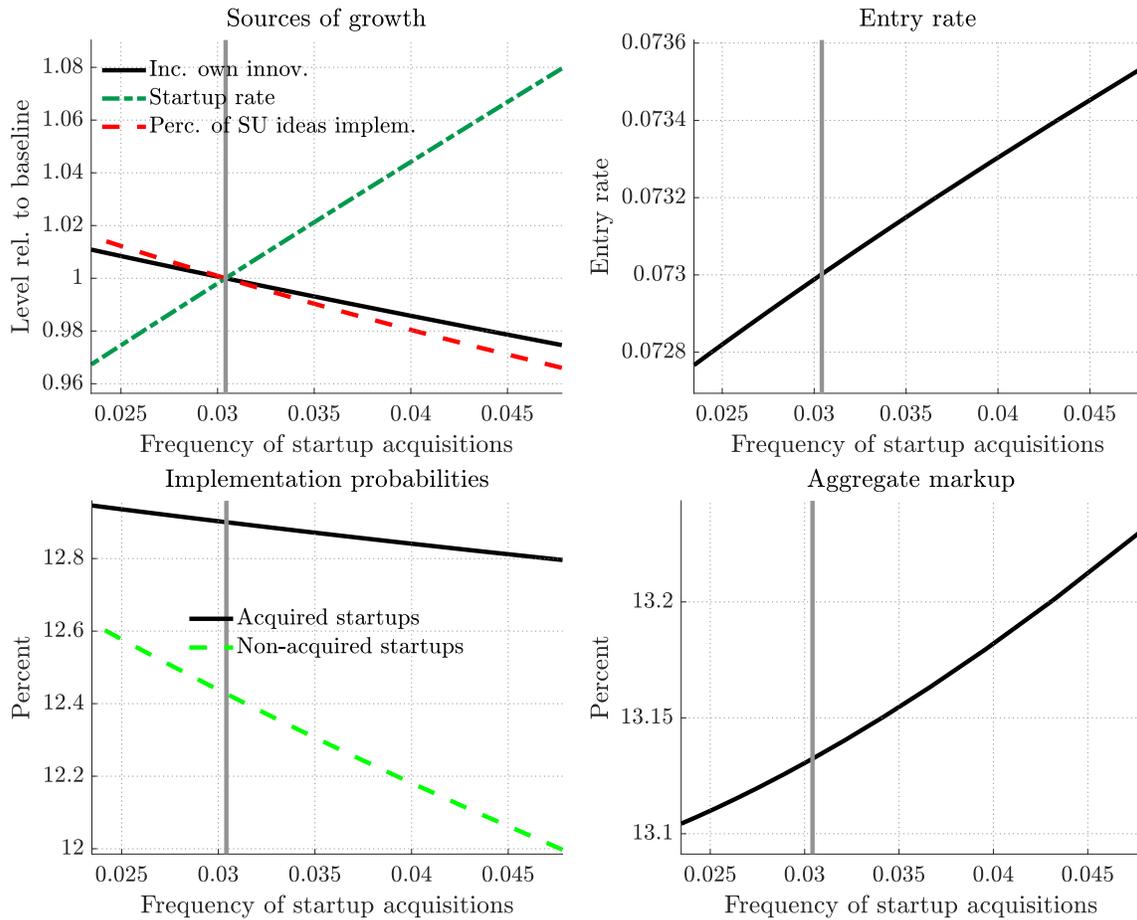


Figure 8: Equilibrium outcomes in selected variables for different values of search costs χ . *Notes:* All plots show the frequency of acquisitions on the x -axis, and all variation in this frequency is driven by changes in search costs. The vertical line marks the baseline frequency of acquisitions.

The decrease in incumbent value due to a higher startup rate explains the decrease in incumbent's own innovation. It also explains that both incumbents and non-acquired startups (which aspire to become incumbents) invest less into idea implementation, as shown in the bottom left panel of Figure 8. This general equilibrium effect, dragging both incumbent and startup implementation down, dominates the partial equilibrium finding that

transferring ideas from startups to incumbents increases their implementation probability.

Finally, more frequent acquisitions also trigger a composition effect. In our baseline calibration, the increase in the startup rate overcompensates the increase in the percentage of acquired startups, so that entry slightly increases (see the top right panel of Figure 8). Together with the fall in incumbent innovation, this shifts the invariant distribution of technology gaps, affecting the aggregate innovation rate. However, this shift is quantitatively very small, in line with the small changes in the entry rate and in the aggregate markup (computed as the inverse of the aggregate labor share) shown in Figure 8.

Summing up, startup acquisitions increase startup creation, but decrease incumbent innovation and implementation by non-acquired startups. In the aggregate, the fall in incumbent innovation turns out to be the dominant force, so that more frequent startup acquisitions are associated with lower growth. Indeed, in our baseline calibration, the contribution of incumbents to growth is $\text{Share}_{\text{inc}} = 71.7\%$.²⁷ Thus, changes in incumbent innovation have an outsize importance for the aggregate growth rate.

The results in this section suggest that policy limits on startup acquisitions could potentially increase growth. The next section examines this issue in greater depth.

4.3 Growth effects of limits to startup acquisitions

We first consider the simplest possible policy: an outright ban on startup acquisitions. This policy might seem radical, but it is actually in line with the Platform Competition and Opportunity Act currently debated by the US Congress.²⁸

Table 7 shows that this policy increases aggregate growth by 0.03 percentage points (or 1.6%) per year. This is the net effect of a 14.9% decrease in the startup rate, which is more than compensated by a 5.3% increase in the own innovation effort of incumbents and a 8.4% increase in the sales-weighted percentage of implemented startup ideas.²⁹

These effects are due to the mechanisms discussed previously. An acquisition ban strongly lowers the incentives for startup creation. However, the fall in startup creation is a boon for incumbents, which now avoid displacement and costly acquisitions, achieve a higher value, and therefore have higher innovation incentives. As incumbent innovation is

²⁷The share of growth due to startup ideas (28%) is close to the share of growth due to entry (26%), which was a calibration target. This is because entry is much more frequent than implementation of a startup idea by an incumbent: few startups are acquired, and acquisitions only slightly raise the implementation probability.

²⁸The Act (aimed at Technology Platforms) would prohibit all acquisitions of direct, nascent or potential competitors, and in our baseline model, all startups are acquired by a competing incumbent. In Section 5.2, we analyse an extended model in which incumbents may also acquire non-competing startups.

²⁹The response of startup creation is close to the one estimated by Eisfeld (2023) for the enterprise software market. Eisfeld estimates that if all startup acquisitions were blocked, startup creation would fall by 8-20%.

the most important source of growth, this effect dominates in the aggregate. Accordingly, the acquisition ban increases consumption-equivalent welfare by 1.8%.³⁰

Table 7: The effects of a startup acquisition ban.

Outcome	Baseline	Acq. Ban	Change
Growth rate	2.00%	2.03%	+1.6%
Incumbent own inn. rate	0.494	0.519	+5.3%
Startup rate	0.760	0.647	-14.9%
Sales-weigh. % of impl. startup ideas	18.1%	19.6%	+8.4%
Entry rate	7.3%	7.2%	-1.6%
Percentage of startups acquired	4.0%	0%	-100%
Aggregate markup	13.1%	13.1%	-0.4%
Consumption-equiv. welfare			+1.8%

Notes: In this table, we compare our baseline BGP to an alternative “acquisition ban” BGP. To compute the latter BGP equilibrium, we impose that the surplus from startup acquisitions is always zero (as it would be, e.g., if a government were to impose an arbitrarily high tax on acquisitions).

Next, Table 8 considers a less radical policy, which only bans startup acquisitions for incumbents with the highest technology gaps. As we had shown earlier (see Figure 4), these acquisitions are most likely to be killer acquisitions. Thus, from a partial equilibrium viewpoint, banning them and maintaining acquisitions that increase the implementation probability of startup ideas appears to be a better policy than a complete ban.

Table 8: The effects of partial acquisition bans.

Change in outcome	Acq. Ban	Ban $n \geq 2$	Ban $n \geq 3$
Growth rate	+1.6%	+1.3%	+0.7%
Incumbent own inn. rate	+5.3%	+4.9%	+3.8%
Startup rate	-14.9%	-14.6%	-13.8%
Sales-weigh. % of impl. startup ideas	+8.4%	+8.3%	+7.8%
Frequency of acquisitions	-100%	-89%	-74%
Consumption-equiv. welfare	+1.8%	+1.6%	+1.0%

Notes: This table compares the effects of various “acquisition ban” BGPs. For these, we impose that the surplus from acquisitions for incumbents with a technology gap exceeding some cutoff is zero (as it would be if a government were to impose an arbitrarily high tax).

³⁰In our model, the representative consumer’s consumption-equivalent welfare change from going from a BGP A to another BGP B is given by $C_0^B \exp\left(\frac{s_B - s_A}{\rho}\right) / C_0^A$.

As Table 8 shows, however, partial bans actually achieve a smaller increase in the growth rate than the complete ban. This is due to general equilibrium effects: a partial ban decreases the startup rate by less (as startups still have some possibilities of selling out), and therefore delivers a smaller boost to incumbent innovation incentives.

In the remainder of the paper, we explore the robustness of our baseline results to different choices for the targeted moments and the externally calibrated parameters. We also consider an extended version of our baseline model, in which we allow incumbents to acquire non-competing startups.

5 Robustness checks and extensions

5.1 Robustness to alternative calibrations

The negative link between startup acquisitions and growth in the baseline calibration is not a hard-wired feature of our model: it is easy to construct alternative calibrations in which our model predicts a positive link.³¹ Thus, our results very much depend on the data used to discipline the model. In this section, we systematically explore the role of different calibration targets and externally calibrated parameters.

Alternative calibration targets To analyze the role of different calibration targets for our results, we re-calibrate our model by changing one calibration target at the time, leaving all other targets and all external parameters unchanged. Figure 9 plots the results for this exercise, using different targets for the causal impact of acquisitions on the number of citations of a startup patent (shown on the x -axis). The baseline value for this target, obtained in our empirical analysis in Section 3.3, was 22%, indicated by the vertical lines in the figure. Each point on the left panel shows the growth effect of an acquisition ban for a re-calibrated model with a different target, and the right panel decomposes this growth effect into the familiar three margins.

Figure 9 shows that, unsurprisingly, targeting a more positive effect of acquisitions on citations means that an acquisition ban provides a smaller boost to growth. As the right panel shows, when acquisitions substantially increase citations, the ban lowers the

³¹For instance, when incumbents have infinite research costs ($\zeta_I \rightarrow +\infty$) and startups have infinite implementation costs ($\kappa_S \rightarrow +\infty$), startups are the only firms that can create ideas, and incumbents are the only firms that can implement them. Then, startup acquisitions are the only way to generate innovations, and prohibiting them would reduce the growth rate to zero. This example is extreme, but it suggests an important insight, formalized in Figure 9: the greater the comparative advantage of incumbents for implementation, the more desirable are startup acquisitions.

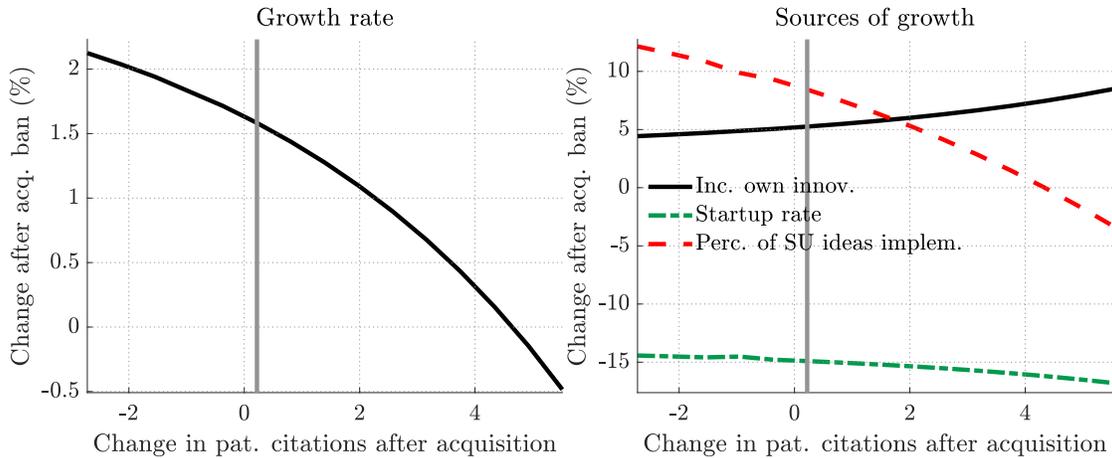


Figure 9: Robustness with respect to different estimates for the causal effect of acquisitions on the number of citations of startup patents. *Notes:* Each point on the x -axis corresponds to a re-calibration of the model, leaving all other calibration targets and all external parameters unchanged. In all calibrations, the model exactly matches all moments. Detailed results for these calibrations are available upon request.

sales-weighted percentage of implemented startup ideas, and this drags the growth rate down. However, Figure 9 also shows that the boost to patent citations would have to be very large in order to reverse our baseline results: a ban lowers growth only if acquisition boosts citations by a factor of 4.2, which is more than 16 times larger than the regression coefficient in our empirical analysis.³² Thus, for all empirically realistic values of our regression coefficient, the negative overall effect of startup acquisitions prevails.

Figure 10 conducts the same exercise by considering different values for the contribution of entrants to growth. As the left panel shows, if this target exceeds its baseline value of 25.7%, acquisition bans become less effective. Indeed, a greater role for entrants implies that startup ideas contribute more to growth. Thus, the decrease in the startup rate after an acquisition ban has a greater impact on aggregate growth. Accordingly, note that the three margins shown in the right panel change relatively little. Changes in aggregate growth are mostly due to changes in the weighting of these margins.

Summing up, our analysis shows that our baseline results are driven by two empirical facts: incumbents account for the bulk of growth, and acquisitions only give a small boost to idea implementation. In environments where entry accounts for a large share of growth and acquisition gives a much larger boost to implementation, startup acquisitions are more benign, and banning them might lower growth. In Appendix C.2, we expand on this analysis by conducting further robustness checks with respect to all other calibration targets.

³²Alternatively, one might think that the elasticity of idea implementation to patent citations is higher than our baseline calibration value. However, Figure 9 shows that it would have to be an order of magnitude higher to overturn our results.

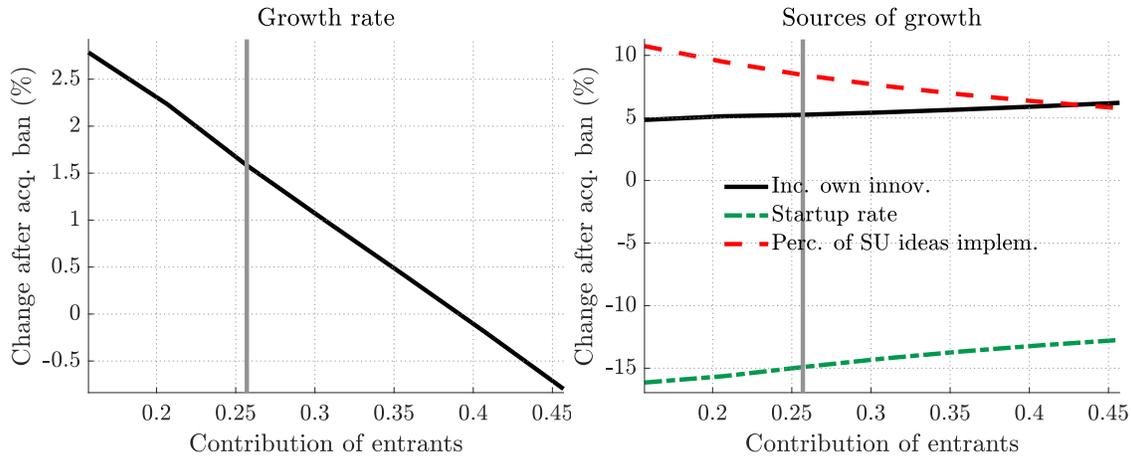


Figure 10: Robustness: the role of the contribution of entrants to growth. *Notes:* See Figure 9.

Alternative values for external parameters Finally, Figure 11 shows the effects of an acquisition ban for different values of incumbent bargaining power α . For each different value of α , we re-calibrate the model to match our baseline targets (leaving all other external parameter values unchanged). The figure shows that greater bargaining power of incumbents reduces the positive effects of an acquisition ban.

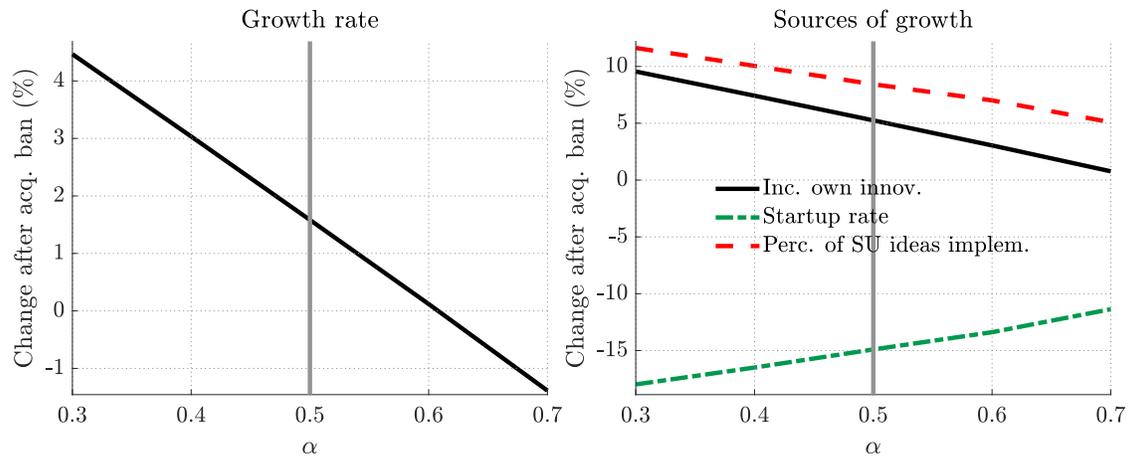


Figure 11: Robustness: the role of incumbent bargaining power. *Notes:* See Figure 9.

Greater bargaining power for incumbents means that the incentive effect of acquisitions on startup creation is weaker. Thus, as shown in the right panel, the startup rate does not fall as much after an acquisition ban. All else equal, this would suggest an even more positive growth effect for an acquisition ban. However, it is more than compensated by incumbent innovation. With high bargaining power, incumbent's value does not suffer much from startup activity, as startup acquisitions are cheap. When α is high enough, a higher

startup rate is even good news for incumbents, as it implies a possibility to cheaply acquire ideas. Accordingly, for high values of α , an acquisition ban does not increase incumbent innovation, and therefore does not increase overall growth.

Overall, the robustness checks in this section suggest that our quantitative results are robust to reasonable variations in our calibration targets and external parameter choices. However, they also suggest that the effects of acquisitions on growth do depend on the characteristics of the economy (and could potentially be different across industries). Precisely, we would expect acquisitions to be more detrimental to innovation when incumbents have a small implementation advantage, are responsible for the bulk of innovation, and have low bargaining power.

5.2 Multiproduct firms and acquisitions of non-competing startups

In our baseline model, incumbents only acquire startups that could potentially displace them. This is an important limitation: in reality, incumbents also buy non-competing startups in order to add new products to their portfolio. In this section, we briefly discuss an extended model that allows for such acquisitions of non-competing startups.

Assumptions We now assume that incumbents can produce multiple products. For each product, incumbents invest into R&D as in the baseline model, and set a search effort s_R for “related” startups (i.e., startups with an idea on the same product, which could potentially displace them). The cost of a search effort s_R for a given product is $\chi_R \cdot (s_R)^\psi \cdot Y_t$.

As in the baseline model, a startup’s idea applies to a randomly chosen product j in the interval $[0, 1]$, and there is a probability $s_R(j)$ that the incumbent producer of j meets the startup. However, we now assume that startups who do not meet the incumbent producer of their product j are matched to another incumbent, currently producing product j' , chosen randomly in the interval $[0, 1]$.³³ When this unrelated incumbent meets the startup, there can be an acquisition, allowing the unrelated incumbent (if it implements the startup’s idea) to expand in the product space by adding product j to its portfolio.

We assume that for an incumbent currently producing k products, the cost of searching for unrelated startups is given by $\chi_U \cdot k \cdot (s_U)^\varphi \cdot Y_t$. By investing this cost, the firm has a probability s_U of meeting an unrelated startup that is matched to any of its products. As in the seminal model of [Klette and Kortum \(2004\)](#), both the cost function and the arrival rate of meetings with unrelated startups are linear in the number of products, implying that all

³³As there is a continuum of products, the probability that j and j' coincide is zero.

firms choose the same search intensity s_U .³⁴

We discuss the solution of this extended model in Appendix D. However, it is worth noting that acquisitions of unrelated startups are not subject to a rent-preserving motive or to the Arrow replacement effect. Therefore, they take place if and only if incumbents have lower implementation costs for ideas, and they always increase the likelihood that the startup’s idea is implemented. This is in line with our empirical evidence in Section 3.4, which showed that acquisitions of unrelated startups lead to a larger increase in the citations of the acquired patents.

Calibration and quantitative results The extended model has one additional parameter with respect to the baseline: the scaling factor for the search cost function for unrelated startups, χ_U . We calibrate this parameter internally and identify it by targeting the share of unrelated startup acquisitions in our sample. We define unrelated startup acquisitions as deals in which the acquirer and the target belong to different SIC 3-digit industries, obtaining a share of 59%.³⁵ The multiproduct model still fits the data very well, as shown in Table D.1 in Appendix D.

Table 9: Acquisition bans in the baseline model and in the multiproduct extension.

Change in outcome	Baseline	Multiproduct		
	Acq. Ban	Acq. Ban	R Acq. Ban	U Acq. Ban
Growth rate	+1.6%	+1.3%	+1.2%	+0.1%
Incumbent own inn. rate	+5.3%	+3.5%	+3.3%	+0.2%
Startup rate	-14.9%	-12.7%	-13.0%	+0.3%
Sales-weigh. % of impl. startup ideas	+8.4%	+9.8%	+10.3%	-0.4%
Frequency of acquisitions	-100%	-100%	-42%	-59%
Consumption-equiv. welfare	+1.8%	+1.7%	+1.6%	+0.1%

Notes: In this table, we compare the effects of different acquisition bans. To compute equilibria with bans, we impose that the surplus from acquisitions is always zero (as it would be if a government were to impose an arbitrarily high tax on acquisitions). For the multiproduct firm model, we also consider separate bans on acquisitions of related and unrelated startups.

Table 9 summarizes the effect of startup acquisition bans in this extended model. The second column shows that a ban of all acquisitions has a slightly smaller growth effect

³⁴Note that when $\chi_U \rightarrow +\infty$, we recover our baseline model.

³⁵This is an intentionally conservative choice. SIC 3-digit industries are relatively narrow, and it is conceivable that firms in different industries compete in the same product market. Choosing a lower target yields results that are even closer to the baseline.

in the multiproduct model, raising the growth rate by 1.3% as opposed to 1.6% in the baseline. This is due the fact that unrelated acquisitions have a smaller effect on innovation incentives and growth than related acquisitions, as the surplus from the former is lower than the surplus from the latter.

Next, we consider two more sophisticated policies, a ban limited to acquisitions of related startups, and a ban limited to acquisitions of unrelated startups. The results, shown in the third and fourth columns of Table 9, confirm that related acquisitions have a much larger growth effect: although they only account for 41% of the total volume of deals, banning them achieves virtually the same growth effect than a complete acquisition ban. Instead, banning unrelated acquisitions has almost no effect. Note that this policy lowers the percentage of implemented startup ideas, which roughly cancels out a small increase in incumbent innovation and in the startup rate.³⁶

Summing up, our results for this extended model suggest that negative growth effects of startup acquisitions are driven by acquisitions of related startups. Thus, a competition authority with limited resources should concentrate its attention on these transactions.

6 Conclusion

In this paper, we assess the effect of startup acquisitions on aggregate growth, using a model that takes into account a large number of potential positive and negative effects of these operations. We discipline the model by calibrating it to micro-level data. Our results indicate that more frequent 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 growth rate by around 0.03 percentage points per year, increasing welfare by 1.8%.

These results are driven by the data that we use to discipline the model. Therefore, they could vary depending on the country and time period considered, and they might be heterogeneous by industry. For instance, we show that startup acquisitions are more beneficial when incumbents have decisive implementation advantages, a low ability to come up with their own ideas, and high bargaining power. Further exploring this heterogeneity is a promising path for future research.

³⁶Banning unrelated acquisitions lowers the outside option of startups when bargaining with related incumbents. Thus, incumbents pay lower acquisition prices, their value increases, and both incumbent innovation and startup creation increase.

References

- 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.
- ARGENTE, D., BASLANDZE, S., HANLEY, D. and MOREIRA, S. (2020). *Patents to Products: Product Innovation and Firm Dynamics*. CEPR Discussion Paper 14692.
- 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.
- ARTS, S., HOU, J. and GOMEZ, J. C. (2021). Natural Language Processing to Identify the Creation and Impact of New Technologies in Patent Text: Code, Data, and New Measures. *Research Policy*, **50** (2), 104144.
- 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 Paper 22750.
- BOURREAU, M., JULLIEN, B. and LEFOUILI, Y. (2018). *Mergers and Demand-Enhancing Innovation*. TSE Working Papers 18-907, Toulouse School of Economics (TSE).
- BRUTTI, Z. and ROJAS, L. E. (2021). M&A and Early Investment Decisions by Digital Platforms. *Mimeo*.
- 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 Paper 13115.
- CALLANDER, S. and MATOUSCHEK, N. (2020). The Novelty of Innovation: Competition, Disruption, and Antitrust Policy. *Working paper*.
- CAO, L. and ZHU, L. (2022). *Mergers and Acquisitions and the Aggregate Markup*. Working paper.
- CAVENAILE, L., CELIK, M. A. and TIAN, X. (2021). The Dynamic Effects of Antitrust Policy on Growth and Welfare. *Journal of Monetary Economics*, **121**, 42–59.

- CELIK, M. A., TIAN, X. and WANG, W. (2022). Acquiring Innovation under Information Frictions. *The Review of Financial Studies*, **35** (10), 4474–4517.
- CUNNINGHAM, C., EDERER, F. and MA, S. (2021). Killer Acquisitions. *Journal of Political Economy*, **129** (3), 649–702.
- DARUICH, D. (2022). The Macroeconomic Consequences of Early Childhood Development Policies. *Mimeo*.
- DAVID, J. M. (2020). The Aggregate Implications of Mergers and Acquisitions. *The Review of Economic Studies*, **88** (4), 1796–1830.
- DECKER, R. A., HALTIWANGER, J., JARMIN, R. S. and MIRANDA, J. (2016). Where Has All the Skewness Gone? The Decline in High-Growth (Young) firms in the U.S. *European Economic Review*, **86**, 4–23.
- 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.
- EISFELD, L. (2023). *Entry and Acquisitions in Software Markets*. Working paper.
- EWENS, M. and FARRE-MENSA, J. (2020). The Deregulation of the Private Equity Markets and the Decline in IPOs. *The Review of Financial Studies*, **33** (12), 5463–5509.
- FEDERICO, G., LANGUS, G. and VALLETTI, T. (2017). A Simple Model of Mergers and Innovation. *Economics Letters*, **157**, 136 – 140.
- FTC (2021). *Non-HSR Reported Acquisitions by Select Technology Platforms, 2010-2019: An FTC Study*. Federal Trade Commission.
- 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.
- GARCIA-MACIA, D., HSIEH, C.-T. and KLENOW, P. J. (2019). How Destructive Is Innovation? *Econometrica*, **87** (5), 1507–1541.
- GUZMAN, J. and STERN, S. (2020). The State of American Entrepreneurship: New Estimates of the Quantity and Quality of Entrepreneurship for 32 US States, 1988–2014. *American Economic Journal: Economic Policy*, **12** (4), 212–43.
- HAUCAP, J., RASCH, A. and STIEBALE, J. (2019). How Mergers Affect Innovation: Theory and Evidence. *International Journal of Industrial Organization*, **63**, 283 – 325.
- HURST, E. and PUGSLEY, B. (2011). What Do Small Businesses Do? *Brookings Papers on Economic Activity*, **42** (2 (Fall)), 73–142.
- 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*. NBER Working Paper 27146.
- 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.
- LANJOUW, J. and SCHANKERMAN, M. (1999). *The Quality of Ideas: Measuring Innovation with Multiple Indicators*. NBER Working Paper 7345.
- 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 Paper 15167.
- LIU, O. (2022). *The Dynamics of Firm Size Inequality: The Role of Acquisition and Innovation*. Working paper.
- PEARCE, J. and WU, L. (2022). Brand Reallocation, Concentration and Growth. *Mimeo*.
- PELLEGRINO, B. (2022). *Product Differentiation and Oligopoly: a Network Approach*. Working paper.
- 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.
- TRAJTENBERG, M., HENDERSON, R. and JAFFE, A. (1997). University versus Corporate Patents: A Window on the Basicness of Invention. *Economics of Innovation and new technology*, **5** (1), 19–50.
- WEISS, J. (2022). *Market Concentration, Growth, and Acquisitions*. Working paper.
- 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 (for online publication only)

A Derivations and proofs

A.1 BGP equilibrium conditions

Using our value function guess, $v(n, \omega) = V_t(\omega, n)/Y_t$, we rewrite equations (9) to (13) as

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

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

$$\sigma(\omega, n) = \max \left[0, \max_{i_A} \left\{ v(\omega, n) + i_A \cdot \left(\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot v(\omega, n + n_S) - v(\omega, n) \right) - \kappa_I \cdot i_A^\psi \right\} - v^{\text{NoMeet}}(\omega, n) - v_S^{\text{NoMeet}}(\omega, n) \right] \quad (\text{A.3})$$

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

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

In all of these expressions, lower-case letters denote values that are normalized by aggregate GDP (e.g., $v^{\text{NoMeet}}(\omega, n) = V_t^{\text{NoMeet}}(\omega, n)/Y_t$, and so on). Using these expressions and the Euler equation (4), we can rewrite the HJB equation (8) as

$$\begin{aligned} \rho \cdot v(\omega, n) = \max_{z, s} \left\{ \omega \cdot \left(1 - \frac{1}{\mu(n)} \right) - \zeta_I \cdot z^\psi - \chi \cdot s^\varphi \right. \\ \left. + z \cdot \max_{i_I} \left[i_I \cdot \left(v(\omega, n + 1) - v(\omega, n) \right) - \kappa_I \cdot i_I^\psi \right] \right. \\ \left. + x \cdot \left[s \cdot \alpha \cdot \sigma(\omega, n) - i_S(\omega, n) \cdot v(\omega, n) \right] \right\} \\ + \sum_{\omega'} \tau_{\omega, \omega'} \cdot \left[v(\omega', n) - v(\omega, n) \right]. \quad (\text{A.6}) \end{aligned}$$

This equation pins down the incumbent value function v as a function of the (endogenous) aggregate startup rate x and the startup's implementation policy function i_S .³⁷ Taking first-order conditions, the incumbent's optimal research investment is

$$z(\omega, n) = \left[\frac{i_I(\omega, n) \cdot (v(\omega, n+1) - v(\omega, n)) - \kappa_I \cdot (i_I(\omega, n))^\psi}{\xi_I \cdot \psi} \right]^{\frac{1}{\psi-1}} \quad (\text{A.7})$$

where $i_I(\omega, n)$ is the optimal implementation probability chosen by the incumbent for its own ideas. The optimal choice equalizes the marginal cost of research to its marginal benefit, which is the arrival of an unimplemented idea. In turn, the optimal startup search investment holds

$$s(\omega, n) = \left(\frac{x \cdot \alpha \cdot \sigma(\omega, n)}{\chi \cdot \varphi} \right)^{\frac{1}{\varphi-1}}. \quad (\text{A.8})$$

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

Note that this equation implies that the incumbent does not invest into search when the surplus from a meeting is zero. Therefore, the probability that a given startup's idea is implemented is always given by $s(\omega, n) \cdot i_A(\omega, n) + (1 - s(\omega, n)) \cdot i_S(\omega, n)$, even when the meeting surplus is zero. We will use this result repeatedly in subsequent equilibrium conditions.

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

$$i_I(\omega, n) = \left(\frac{v(\omega, n+1) - v(\omega, n)}{\kappa_I \cdot \psi} \right)^{\frac{1}{\psi-1}}. \quad (\text{A.9})$$

Again, firms equalize the marginal cost of implementation to its marginal benefit, which comes from increasing the technology gap. Investment of incumbents into acquired startup ideas holds

$$i_A(\omega, n) = \left(\frac{\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot v(\omega, n+n_S) - v(\omega, n)}{\kappa_I \cdot \psi} \right)^{\frac{1}{\psi-1}} \quad (\text{A.10})$$

Finally, as equation (A.2) shows, the optimal implementation probability chosen by

³⁷Each startup has a Poisson arrival rate 1 of ideas. Thus, x is equal to the mass of startups (both in absolute terms and relative to the mass of incumbents) and the arrival rate of startup ideas.

startups for their own ideas is independent of the technology gap n . It is given by

$$i_S(\omega) = \left(\frac{\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot v(\omega, n_S)}{\kappa_S \cdot \psi} \right)^{\frac{1}{\psi-1}} \quad (\text{A.11})$$

Further, we can rewrite the value of a startup (the right-hand side of equation (14)) as

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

Therefore, the free entry condition becomes

$$\zeta_S = \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega, n) \cdot \left[v_S^{\text{NoMeet}}(\omega) + s(\omega, n) \cdot (1 - \alpha) \cdot \sigma(\omega, n) \right]. \quad (\text{A.12})$$

Creating a startup always delivers the startup outside option v_S^{NoMeet} . Moreover, if the startup meets the incumbent owning the product to which its idea applies, it is potentially acquired and captures a share of the acquisition surplus. Thus, everything else equal, more frequent acquisitions always increase the value of creating a startup.

A.2 The joint distribution of quality and technology gaps

To compute the endogenous joint distribution of products over quality and technology gaps, we build an infinitesimal generator matrix. For a homogeneous continuous-time Markov chain z_t taking values in $\{z_1, z_2, \dots, z_S\} \in \mathbb{R}^S$, the generator matrix \mathbf{M} is:

$$\mathbf{M} \equiv \begin{pmatrix} -\sum_{j \neq 1} \lambda_{1j} & \lambda_{12} & \dots & \lambda_{1S} \\ \lambda_{21} & -\sum_{j \neq 2} \lambda_{2j} & \dots & \lambda_{2S} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{S1} & \lambda_{S2} & \dots & -\sum_{j \neq S} \lambda_{Sj} \end{pmatrix} \quad (\text{A.13})$$

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

To build this matrix in our model, we assume that the technology gap is bounded above

by $n_{max} = 20$, so that the state space is $\Omega \times \{1, 2, \dots, 20\}$.³⁸ Denoting by $m_t(\omega, n)$ the share of firms in state (ω, n) at time t , the law of motion of m_t is:

$$\frac{\partial \vec{m}_t}{\partial t} = \mathbf{M}^\top \vec{m}_t \quad (\text{A.14})$$

where \vec{m}_t is a stacked vector across all states. To find the invariant distribution, we impose $\frac{\partial \vec{m}_t}{\partial t} = \vec{0}$ in equation (A.14) and solve for the unique solution of the system of linear equations holding $\sum_\omega \sum_n m(\omega, n) = 1$.

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

$$z(\omega, n) \cdot i_I(\omega, n) + x \cdot \left(s(\omega, n) \cdot i_A(\omega, n) \cdot \theta(1) + (1 - s(\omega, n)) \cdot i_S(\omega) \cdot \theta(n + 1) \right).$$

Such transitions occur because of incumbent's own innovations, 1-step startup ideas implemented by incumbents, and $(n + 1)$ -step startup ideas implemented by startups.

For transitions from $(\omega, n_{max} - 1)$ to (ω, n_{max}) , the transition rate is

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

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 (ω, n) to $(\omega, n + k)$, with $k > 1$ and $n + k < n_{max}$, we have a transition rate

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

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 (ω, n) to (ω, n_{max}) , with $n < n_{max} - 1$, we have

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

³⁸We verify that the mass of products that have a technology gap equal to 20 is always very small. Therefore, increasing n_{max} to 25 or 30 does not affect our results.

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.

For downward transitions, from (ω, n) to (ω, k) with $n > k$, we have a transition rate

$$x \cdot (1 - s(\omega, n)) \cdot i_S(\omega) \cdot \theta(k)$$

Downward transition occur only through entry, when the entering startup takes k steps.

Finally, transitions from (ω, n) to (ω', n) are exogenous, and occur at rate $\tau_{\omega, \omega'}$. Knowing the full generator matrix, we can solve for the invariant distribution of products over (ω, n) .

Summing up, given innovation, implementation, search and acquisition decisions, we can solve for the invariant distribution of products over quality and technology gaps. These decisions, however, in turn depend on the distribution. In the next section, we discuss the algorithm that we use to jointly solve for all these outcomes.

A.3 Solution algorithm

To solve for the startup rate, innovation, search and acquisition choices, and the invariant distribution of products across quality and technology gaps, we use the following algorithm.

1. We guess a value for the startup rate x .
2. Given this guess, we solve for the value function of incumbent firms, using the following value function iteration algorithm.
 - (a) We guess a value function v .
 - (b) Using equations (A.7) to (A.11), we deduce from this guess the policy functions z, s, i_L, i_A, i_S as well as the acquisition surplus σ .
 - (c) We use equation (A.6) to compute a new implied value for the value function, v_{new} .
 - (d) If $\| \frac{v - v_{new}}{v_{new}} \|_{\infty} < 10^{-4}$, the algorithm has converged and we proceed to step 3. If this condition does not hold, we compute a new guess for the value function as $0.998 \cdot v + 0.002 \cdot v_{new}$ and go back to step 2 (b).
3. Using the innovation rates obtained in step 2 and our guess for the startup rate x , we compute the joint distribution of products across quality levels and technology gaps, as described in Appendix A.2.

4. We compute the value of creating a startup v_S (the right-hand side of equation (A.12)), using our result for the value function of incumbents obtained in step 2 and the distribution of incumbents obtained in step 3. When the condition

$$\left| \frac{\bar{\xi}_S - v_S}{v_S} \right| < 10^{-4}$$

holds, we have found the equilibrium. Otherwise, we update our guess for x and return to step 2.

Note that the algorithm is independent of the level of other aggregate variables such as the wage w_t or the growth rate g . Thus, we can solve for these outcomes after computing firm-level decisions. This block structure greatly simplifies the solution of the model.

A.4 Growth and other aggregate variables

Product market clearing implies $Y_t = C_t + R_t + I_t + S_t$, where:

$$R_t = Y_t \cdot \left(x \cdot \bar{\xi}_S + \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega, n) \cdot \bar{\xi}_I \cdot (z(\omega, n))^\psi \right) \quad (\text{A.15})$$

$$I_t = Y_t \cdot \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega, n) \cdot \left[z(\omega, n) \cdot \kappa_I \cdot (i_I(\omega, n))^\psi + x \cdot \left(s(\omega, n) \cdot \kappa_I \cdot (i_A(\omega, n))^\psi + (1 - s(\omega, n)) \cdot \kappa_S \cdot (i_S(\omega))^\psi \right) \right], \quad (\text{A.16})$$

$$S_t = Y_t \cdot \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega, n) \cdot \chi \cdot (s(\omega, n))^\varphi, \quad (\text{A.17})$$

This shows that consumption is always proportional to GDP, and thus grows at the same rate. We next derive this common growth rate.

In line with the literature, we define the aggregate markup μ as the inverse of the labor share (i.e., $\mu \equiv \frac{Y_t}{w_t \cdot L}$). This aggregate markup is pinned down by equation (15). It is then easy to show that the labor used to produce product j at time t holds $l_{jt} = \omega_{jt} \cdot \frac{\mu}{\mu(n_{jt})} \cdot L$. Using this result, we can show

$$Y_t = A_t \cdot \mathcal{M} \cdot L, \quad (\text{A.18})$$

where A_t is aggregate productivity, defined as

$$A_t \equiv \exp \left(\int_0^1 \omega_{jt} \cdot \ln(a_{jt}) \, dj \right). \quad (\text{A.19})$$

and \mathcal{M} is given by

$$\mathcal{M} = \exp \left(\sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega, n) \cdot \omega \cdot \ln \left(\frac{\mu}{\mu(n)} \right) \right).$$

The \mathcal{M} term captures the static misallocation of labor due to markup dispersion, as in [Peters \(2020\)](#). When all firms charge the same markups, $\mathcal{M} = 1$. Crucially, the above expression shows that \mathcal{M} is constant on the BGP. Therefore, output grows at the same rate as aggregate productivity A_t .

To derive the growth rate of aggregate productivity, we write

$$\begin{aligned} \ln(A_t) &= \int_0^1 \omega_{j_t} \cdot \ln(a_{j_t}) dj \\ &= \sum_{a \in \mathbb{A}} \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m_t(\omega, n, a) \cdot \omega \cdot \ln(a), \end{aligned}$$

where $m_t(\omega, n, a)$ stands for the mass of goods with quality ω , technology gap n and productivity a at instant t . Note that because productivity moves on a ladder, it can only take values in a countable set \mathbb{A} .

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 aggregate productivity at instant $t + dt$ as

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

Rearranging the sum, and using that by definition, $\sum_{a \in \mathbb{A}} m_t(\omega, n, a) = m(\omega, n)$, we get

$$\frac{\ln(A_{t+dt}) - \ln(A_t)}{dt} = \ln(\lambda) \cdot \sum_{\omega \in \Omega} \left(\sum_{n=1}^{+\infty} m(\omega, n) \cdot \omega \cdot \sum_{k=1}^{+\infty} k \cdot b(\omega, n, k) \right). \quad (\text{A.20})$$

Taking the limit as dt goes to 0 yields the growth rate of aggregate productivity:

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

As a final step, we can further simplify the inner sum. To do so, we note

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

The last term in square brackets is simply the expected value of the steps taken by a successful startup innovation, and therefore equal to $1 + \gamma$. Putting these results together, we obtain equation (17) in the main text.

A.5 Decomposition formula

We can write the growth formula (17) as

$$g = \ln(\lambda) \cdot \left(\mathcal{I} + (1 + \gamma) \cdot x \cdot \mathcal{P} \right), \quad (\text{A.21})$$

where \mathcal{I} and \mathcal{P} are the sales-weighted averages of incumbent innovation and of the implementation probability of startup ideas, as defined in equation (17), that is:

$$\mathcal{I} \equiv \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega, n) \cdot \omega \cdot z(\omega, n) \cdot i_I(\omega, n) \quad (\text{A.22})$$

$$\mathcal{P} \equiv \sum_{\omega \in \Omega} \sum_{n=1}^{+\infty} m(\omega, n) \cdot \omega \cdot \left(s(\omega, n) \cdot i_A(\omega, n) + (1 - s(\omega, n)) \cdot i_S(\omega, n) \right) \quad (\text{A.23})$$

To get the decomposition, we can simply rewrite (A.21) as:

$$\frac{g^B}{g^A} = \frac{\ln(\lambda) \cdot \mathcal{I}^A}{g^A} \cdot \frac{\mathcal{I}^B}{\mathcal{I}^A} + \frac{\ln(\lambda) (1 + \gamma) \cdot x^A \cdot \mathcal{P}^A}{g^A} \cdot \frac{x^B}{x^A} \cdot \frac{\mathcal{P}^B}{\mathcal{P}^A}, \quad (\text{A.24})$$

which is equation (19), with $\text{Share}_{\text{inc}}^A \equiv \frac{\ln(\lambda) \cdot \mathcal{I}^A}{g^A}$.

B Data appendix

B.1 Data construction

Our data sources are fully described in the main text. In order to match the firms in the acquisitions database to their patents and accounts, we follow different routes for publicly listed and private firms.

For publicly listed firms, we use the fact that ThomsonONE provides us with a firm identifier, GVKEY, which is also used in the NBER Patent Data Project database and in Compustat. However, GVKEY is not a permanent firm identifier, as may change for instance after major mergers and acquisitions. Therefore, we first need to match each GVKEY to a permanent firm identifier, PERMCO, provided by the Center for Research in Security Prices (CRSP), and then match all data sets based on this PERMCO.

For private firms, we cannot follow this procedure, as there are no common identifiers across datasets. Instead, we match these firms based on their names. To do so, we first standardize names, by removing commonly used endings such as "CO", "CORP" or "CORPORATION". This yields some exact matches, but their number is still relatively small. Therefore, we complement this by a manual matching exercise, considering all private firm names in the acquisition database and checking whether there is a corresponding firm name in the patent database. This catches many cases with minor variations in the spelling of firm names, such as for instance "ACCENT OPTICAL TECH INC" and "ACCENT OPTICAL TECHN INC".

B.2 Matching treated to control patents

This section provides further details on how we match treated to control patents.

We start with a broad sample of USPTO patents for which one of the inventors is located in the US. Next, we identify patents granted to any startup firm that will be acquired in the 6 years following its first patent application. Finally, conditional on this sample, we keep the subset of patents granted at least two years prior to the acquisition, as this allows us to track citations received before the acquisition. This leaves us with 3,040 patents, the *treated patent sample*. For each of these treated patents, we aim to find a set of control patents that looks as similar as possible in terms of observables at the time of the acquisition.

First matching stage In the first stage we do an exact match on two categories: application year (24 values ranging from 1980 to 2003) and technological subcategory (31 values).

In other words, we limit the pool of potential control patents to the ones that have the same application year and the same technological subcategory as the treated patent.

Second matching stage Conditional on satisfying the criteria of the first matching stage, we estimate a propensity score matching algorithm based on several variables, aiming to capture ex-ante intrinsic patent features as well as citations received before acquisition. Variables capturing ex-ante text-based patent features are from [Arts, Hou and Gomez \(2021\)](#). The authors develop natural language processing techniques to identify the creation and impact of new technologies in USPTO utility patents between 1969 and 2018. They validate the new techniques and their improvement over traditional metrics and provide open access to code and data.

We merge their dataset to ours and use the following ex-ante text-based measures: (1) *new word combinations* - contains the number of new pairwise keyword combinations introduced by the patent; (2) *new bigrams* - contains the number of new bigrams (two consecutive keywords in the patent document) introduced by the patent for the first time; (3) *new trigrams* - contains the number of new trigrams (three consecutive keywords in the patent document) introduced by the patent; (4) *novelty* - defined as 1 minus the backward cosine, where the backward cosine contains the average cosine similarity between the focal patent and all other patents filed in the five years before the focal patent;³⁹ (5) *impact* - defined as forward cosine textual similarity divided by backward cosine textual similarity. The intuition is that a patent has a greater impact if it is dissimilar to previous patents, but at the same time many future patents follow on its footsteps.

Moreover, we also match on the following four measures: (6) *originality*, based on [Trajtenberg, Henderson and Jaffe \(1997\)](#), if a patent cites previous patents that belong to a narrow set of technologies the originality score will be low, whereas citing patents in a wide range of fields would render a high score; (7) *number of claims* - claims specify in detail the building blocks of the patented invention, and hence their number is indicative of the scope or width of the invention ([Lanjouw and Schankerman, 1999](#)); (8) citations received in the first year after the patent is granted; (9) citations received in the second year after is granted. We perform the propensity score matching on a caliper width of 0.3 and allow for up to 10 control patents for each treated patent.

³⁹Novelty should capture how different the patent is compared to patents filed in the previous five years. More novel patents are arguably more dissimilar in content compared to prior patents.

B.3 Summary statistics

Table B.1 shows some descriptive statistics on the variables used in our empirical analysis. The dependent variable in our regressions is the number of citations received at the patent-year level, and has a mean of 1.54 in our sample of 206,432 observations. Around 40% of observations satisfy the condition that both the acquirer and the target have the same primary 3-digit SIC code. When decomposing acquirers by industries, around 23% are in the IT sector, 22% in the medical sector, and 19% in the semiconductors sector.

Table B.1: Descriptive statistics

	Obs	Mean	Min	Median	Max
Number of Cites Received	206,432	1.54 (2.28)	0	1	8
Dummy Treatment	206,432	0.09 (0.29)	0	0	1
Dummy Post	206,432	0.62 (0.48)	0	1	1
Acquirer Market Share	206,432	0.01 (0.04)	0	0.01	0.33
Dummy Same SIC3	206,432	0.40 (0.49)	0	0	1
Dummy IT Sector	206,432	0.23 (0.42)	0	0	1
Dummy Medical Devices Sector	206,432	0.22 (0.41)	0	0	1
Dummy Semiconductors Sector	206,432	0.19 (0.14)	0	0	1
Dummy Other Sector	206,432	0.53 (0.50)	0	1	1

Notes: *Number of Cites Received* summarises the values of the (winsorised) dependent variable. *Dummy Treatment* takes value 1 for patents of a target firm that is acquired in the first 6 years since starting to patent. *Dummy Post* takes value 1 for the years after an acquisition. *Acquirer Market Share* is the market share in the primary industry code of the acquiring firm at the time of acquisition. *Dummy Same SIC3* is a dummy taking value 1 when both the acquirer and target firms have the same primary industry code. The remaining variables are dummies based on the sector to which the acquiring firm belongs to.

Table B.2 compares treated patents to control patents after our propensity score matching. The resulting sample of 2,519 treated patents and 25,135 control patents has similar

mean values in all nine variables that were involved in the matching exercise. Regarding the five text-based measures (new word combinations; new bigrams; new trigrams; novelty; impact), the t -test is never statistically significant. Treated and control patents also have similar means in terms of originality and cites received in the initial years. The only statistically significant difference (at the 10% level) is in the number of claims, where treated patents have a slightly higher value. Overall, this table supports the claim that both sets of patents are similar across a large number of observables.

Table B.2: Comparison of matching variables - Treatment vs Control group

	Treatment Patents		Control Patents		t -test
	Obs.	Mean (St.dev.)	Obs.	Mean (St.dev.)	p -value
New Word Combination	2,519	140.81 (415.44)	25,135	131.86 (683.06)	0.52
New Bigrams	2,519	3.15 (4.73)	25,135	3.05 (5.37)	0.35
New Trigrams	2,519	4.97 (6.62)	25,135	4.81 (8.03)	0.35
Novelty	2,519	0.97 (0.01)	25,135	0.97 (0.01)	0.57
Impact	2,519	1.03 (0.15)	25,135	1.03 (0.15)	0.74
Originality	2,519	0.54 (0.31)	25,135	0.54 (0.32)	0.73
Number of Claims	2,519	22.10 (17.56)	25,135	21.38 (18.95)	0.07*
Cites Received 1st Year	2,519	1.71 (3.90)	25,135	1.63 (4.27)	0.37
Cites Received 2nd Year	2,519	4.02 (7.11)	25,135	3.87 (7.81)	0.35

Notes: This table compares treated to control patents for the variables incorporated in the propensity score matching exercise. All variables are defined in Section B.2.

B.4 Additional tables and figures

Table B.3 shows that our baseline findings in Section 3.3 are robust to many alternative specifications. All columns in this robustness table include both year fixed effects and

matched patent pair fixed effects. For comparison purposes, column (1) replicates the main finding in column (4) of Table 3: treated patents receive 21.8% more citations in the post-acquisition period compared to control patents of the same matched patent pair. Column (2) reduces the number of closest control patents for each treated patent from 10 to 5. Column (3) reduces the event study window to 5 years before and after the acquisition, while our baseline is 7. Column (4) does not winsorize the dependent variable. Column (5) includes industry-year fixed effects instead of the baseline year fixed effects. In all these robustness tests using the Poisson estimator, the estimated coefficient stays both economically and statistically significant. Finally, in the last two columns we present results from OLS estimations. In column (6) the dependent variable is in levels, so that the estimated coefficient of the interaction term cannot be interpreted as an elasticity anymore; instead, we find that treated patents in the post-acquisition period will on average receive 0.419 additional citations and the result is highly statistically significant. In column (7), the dependent variable is the log of the number of citations received plus one. The economic magnitude is now slightly smaller than in the Poisson estimations (13.1% vs 21.8%) but still statistically significant at the 1% level.

Appendix Table B.4 splits our sample according to the industry of the acquirer. For acquirers from the IT, semiconductor and medical devices industry, we find always a positive and statistically significant coefficient on the interaction term. The magnitude of these sector-specific coefficients do not differ too much from the value for the full sample (21.8%).

Appendix Table B.5 provides some descriptive statistics on the growth outcomes of startups coming from Guzman and Stern (2020). In column (1) we include all startups (independently of whether they patent or not) and look at some success measures after 6 years of existence. For this full sample we find that 0.23% reach a threshold level of 100 employees or above; 0.06% will be acquired during these initial years; finally, only 0.01% of startups go public in the first 6 years of existence. Column (2) repeats the exercise for the selected subset of startups who patent. The sample shrinks from over 18 million in column (1) to 37,588 startups in column (2). As expected, the probability of success outcomes rises substantially. This subset of firms will reach 100 or more employees with a probability of 6.60%. Also, 4% will be acquired and over 1% will even manage to go public in the initial 6 years of existence. Finally, column (3) chooses an even more selected subset of startups; it only keeps that ones that patent and additionally file in Delaware. Out of these 10,804 startups, almost 14% succeed in having over 100 employees, almost 10% will be acquired, and almost 3% of them will go public.

Table B.3: The effect of acquisitions on the implementation of ideas - Robustness

Dependent variable: number of citations received

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
D(Post)	0.346*** (0.019)	0.343*** (0.020)	0.226*** (0.017)	0.392*** (0.023)	0.373*** (0.018)	0.541*** (0.021)	0.197*** (0.006)
D(Treatment)	-0.010 (0.035)	-0.006 (0.034)	-0.007 (0.034)	-0.010 (0.045)	-0.009 (0.029)	-0.014 (0.039)	-0.006 (0.012)
D(Post) * D(Treatment)	0.218*** (0.041)	0.223*** (0.041)	0.216*** (0.042)	0.263*** (0.061)	0.215*** (0.032)	0.419*** (0.083)	0.131*** (0.025)
Observations	206,352	112,553	186,184	206,352	205,601	206,432	206,432
R-squared	0.297	0.287
Matched Pair FE	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓
Closest control patents		✓					
Pre & Post Year Window			✓				
No winsorization				✓			
Industry × Year FE					✓		
OLS: levels						✓	
OLS: logs							✓

Notes: We use a Poisson estimator. The dependent variable is the number of citations received at the patent-year level. $D(\text{Treatment})_i$ takes value 1 for treated patents, and $D(\text{Post})_{it}$ takes value 1 for the years after the acquisition. For ease of comparison, column (1) repeats our baseline result. Column (2) limits the sample to the 5 closest nearest neighbors for each treated patent. Column (3) shrinks the window of years to 5 years prior acquisition until 5 years post acquisition. Column (4) does not winsorize the dependent variable. Column (5) includes industry-year fixed effects instead of year fixed effects. Column (6) is an OLS regression in levels. Finally, column (7) is an OLS regression in logs. Standard errors are clustered at the target firm level. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table B.4: Industry-level regression results

Dependent variable: number of citations received

	IT	Medical Devices	Semi- conductors	Other Categories
D(Post)	0.214*** (0.022)	0.460*** (0.019)	0.116 (0.099)	0.336*** (0.019)
D(Treatment)	0.033 (0.046)	-0.004 (0.059)	-0.014 (0.181)	0.010 (0.044)
D(Post) * D(Treatment)	0.177*** (0.052)	0.268*** (0.057)	0.335* (0.187)	0.254*** (0.049)
Observations	43,970	46,474	2,796	120,893
Matched Pair FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓

Notes: We use a Poisson estimator. The dependent variable is the number of citations received at the patent-year level. $D(\text{Treatment})_i$ takes value 1 for treated patents, and $D(\text{Post})_{it}$ takes value 1 for the years after the acquisition. Standard errors are clustered at the target firm level. * significant at 10%; ** significant at 5%; *** significant at 1%.

Table B.5: Growth outcomes for newly incorporated firms: more details.

Sample	(1) All new firms	(2) Patenting new firms	(3) Patent & Delaware
Total number	18,764,856	37,588	10,804
<i>Outcome after 6 years</i>			
... Acquisition	0.06%	4.02%	9.32%
... IPO	0.01%	1.13%	2.94%
... 100+ employees	0.23%	6.60%	13.74%

Source: [Guzman and Stern \(2020\)](#) and own computations. The sample contains all newly incorporated firms incorporated in 32 US states between 1988 and 2008. Column (1) refers to all new firms, column (2) to new firms with a patent application, and column (3) to new firms with a patent application and an incorporation in Delaware.

C Calibration details

C.1 Global identification

This section lays out a rigorous test of global identification for the calibration exercise presented in Section 4.1. We use a method that was first developed by [Daruich \(2022\)](#).

The [Daruich \(2022\)](#) method Denote our vector of parameters by $\theta \in \Theta \subset \mathbb{R}_+^M$, where Θ is the parameter space and $M \in \mathbb{N}$ is the number of internally calibrated parameters. First, we create a large M -dimensional hyper-cube $\mathcal{P} \subset \Theta$ in the parameter space (which requires choosing lower and upper bounds on each parameter). Then, we iteratively pick quasi-random realizations from \mathcal{P} using a Sobol sequence, which successively forms finer uniform partitions of the space. For a sufficiently large number of Sobol draws, this routine efficiently and comprehensively explores every corner of \mathcal{P} . For each parameter draw, we 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 thousands of separate CPUs, thereby saving us a great amount of computation time. After N Sobol draws (in practice, $N \approx 2$ million), we have a $N \times M$ matrix \mathbf{R} of results and a $N \times M$ matrix $\mathbf{P} \in \mathcal{P}$ of the corresponding parameters. The model-generated data contained in the (\mathbf{R}, \mathbf{P}) matrices can then be exploited to obtain information about identification.

We implement the following procedure. First, for each parameter $p \in \theta$, 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. With this, we may 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 the following four criteria are satisfied:

- (i) the distribution changes monotonically across quantiles of p ;
- (ii) the rate of this change is high;
- (iii) the inter-quartile range of the m distribution is small throughout the support of p ;
- (iv) at the calibrated value, the empirical target falls within the inter-quartile range.

Criterion (i) implies that m is sensitive to variation in p ; criterion (ii) gives an idea of how strong this sensitivity is; criterion (iii) implies that other parameters are relatively unimportant to explain the moment; and criterion (iv) implies that the empirical target is not an outlier occurrence at the calibrated value of the parameter. Importantly, all the remaining parameters are not fixed throughout this analysis, but vary in a random fashion. Therefore, this method gives us a global view of identification. In doing so, it outperforms identification methods based on local elasticities (i.e., methods based on pseudo-derivatives of moments obtained by keeping all but one of the parameters fixed at their calibrated values).

Identification results Figure C.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 6 and in our verbal discussion in Section 4.1). All in all, we find that all the parameters of the model are very well-identified by all four criteria.

C.2 Additional robustness exercises

Using the methodology introduced in Section 5, Figures C.2 and C.3 show additional robustness results for the growth effects of an acquisition ban, for different re-calibrations of the model across alternative calibration targets. Each row in these figures considers a different targeted moment (the growth rate, the exit rate, the relative size of acquiring firms, the percentage of acquired startups, and the implementation probability of startup ideas). The left column shows the overall growth effect of the ban across the different calibration targets, and the right column shows the usual decomposition into the three sources of growth. Vertical lines indicate the value targeted in the baseline calibration.

Figure C.2 shows that the growth effect of the acquisition ban becomes stronger when we target a lower growth rate, a higher exit rate, a larger relative size of acquiring firms and a higher percentage of acquired startups.

The implementation probability of startup ideas Figure C.3 illustrates the claim made in Section 4.1: different targets for the implementation probability of startup ideas lead to the same quantitative results. In particular, a ban on startup acquisition raises the growth rate by the same amount irrespective of whether the target for the implementation probability is 10, 20 or 25%.

This irrelevance is not a coincidence. First, the target for the implementation rate of

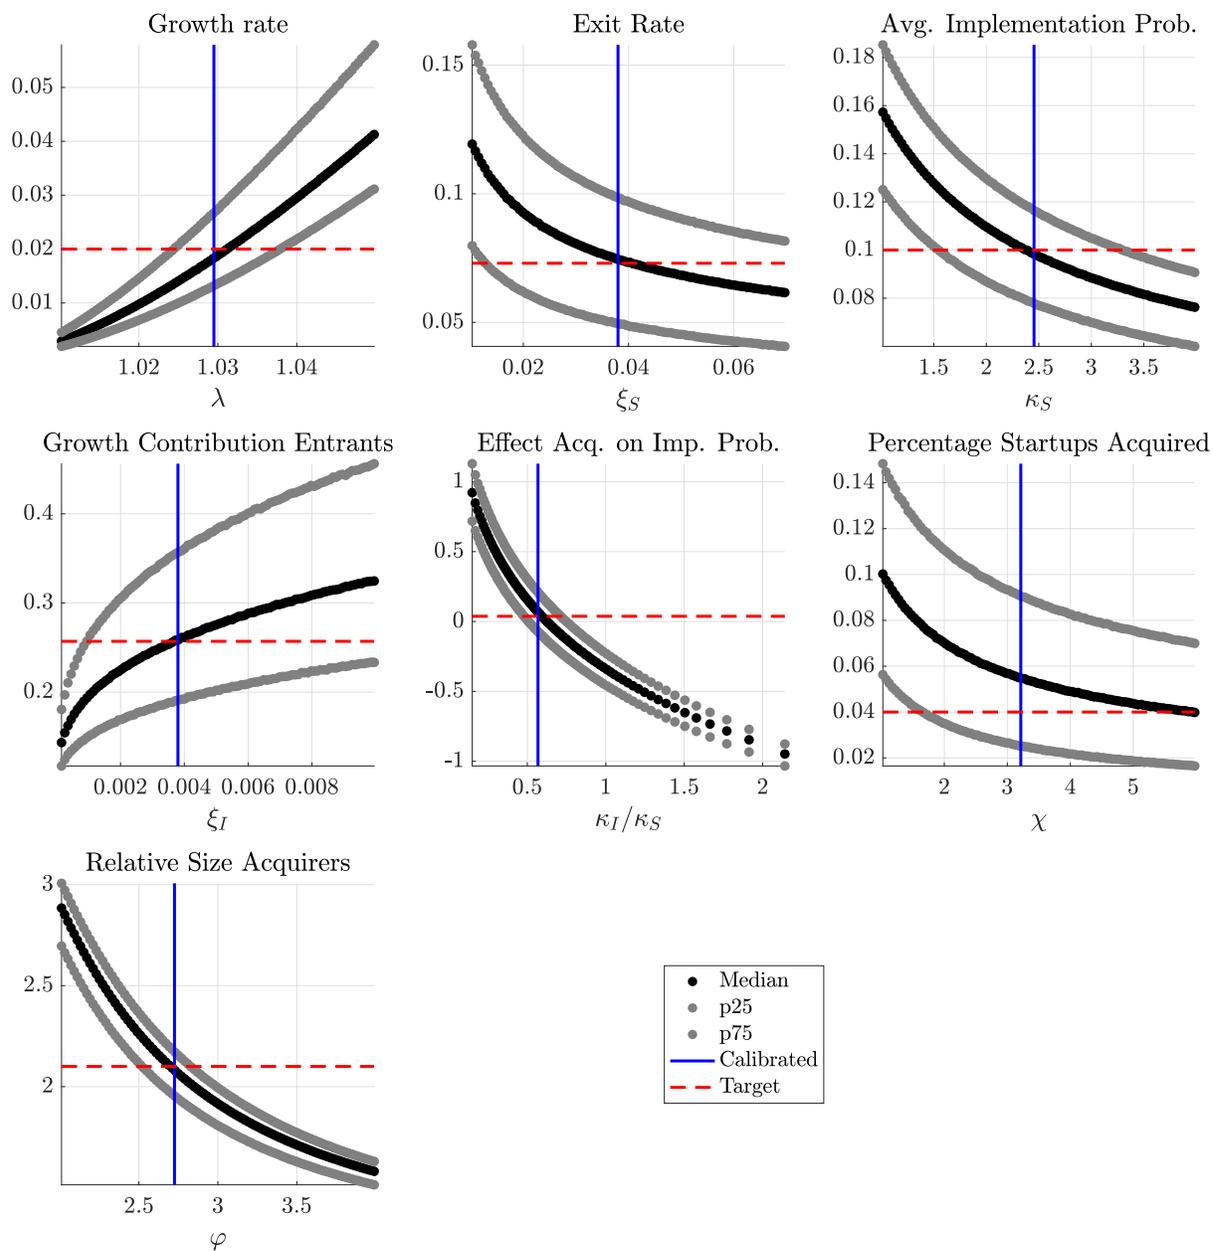


Figure C.1: Global identification results. For each parameter-moment pairing, the black dots are the median of the distribution generated by random variation in all the other parameters, and the gray dots are the inter-quartile range of the distribution. The dashed horizontal line marks the empirical target, and the vertical line marks the calibrated parameter value.

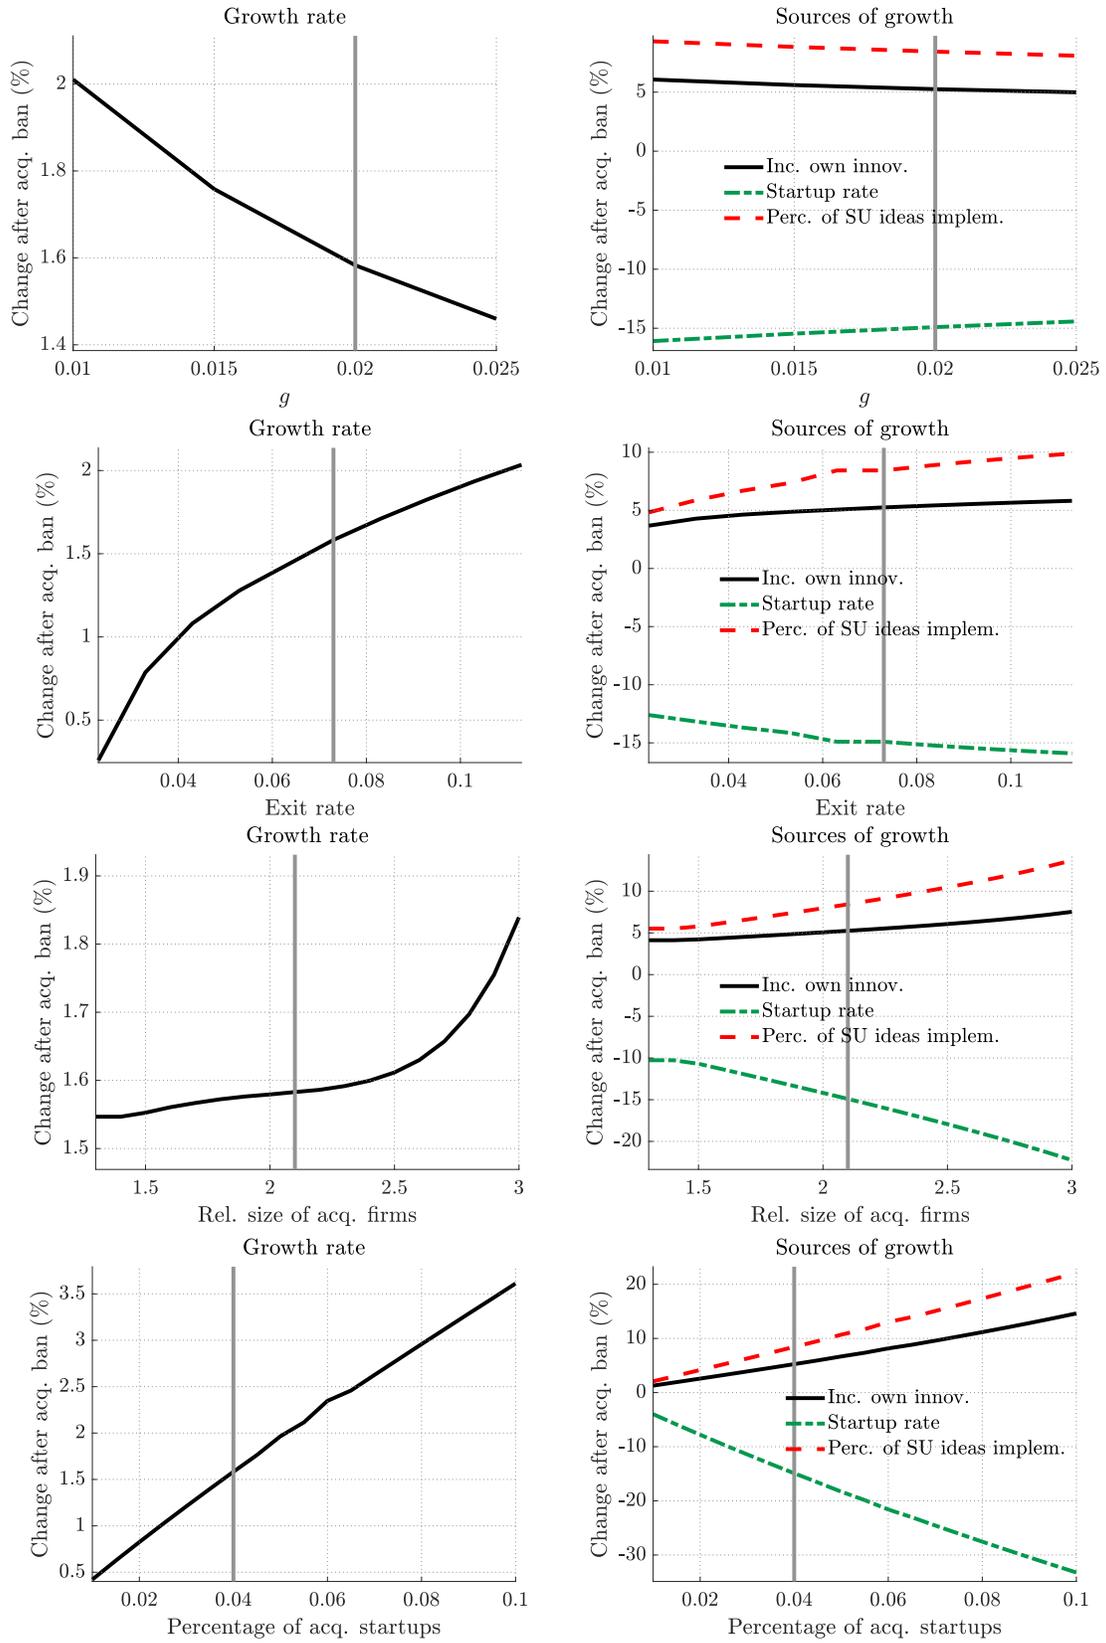


Figure C.2: The role of different calibration targets for the growth effect of a startup acquisition ban. *Notes:* See Figure 9.

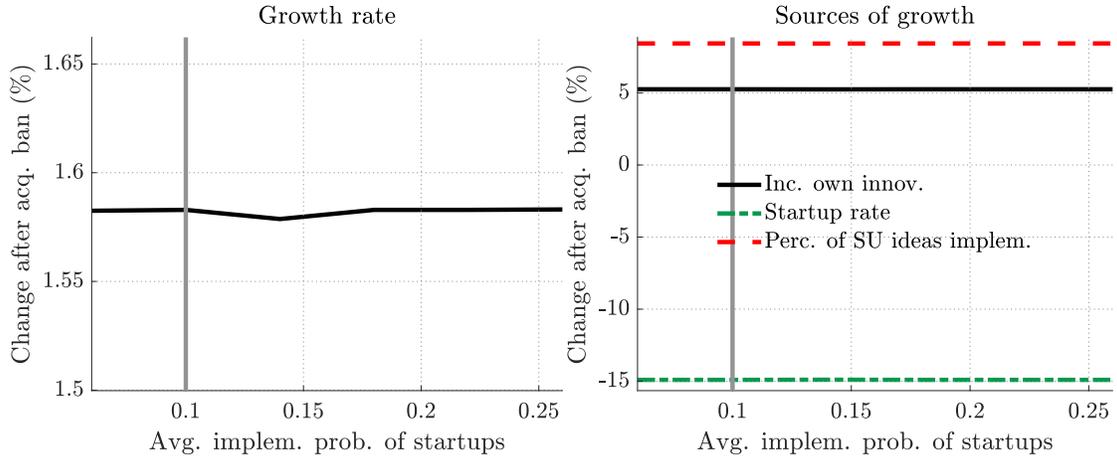


Figure C.3: Robustness: the role of the average implementation probability of startups. *Notes:* See Figure 9.

startups pins down the parameter κ_S , but this parameter is irrelevant for policy analysis in our model. Indeed, κ_S only appears once in the equilibrium conditions pinning down the growth rate, in equation (A.11), determining the implementation choice i_S of non-acquired startups. However, from this equation, it is clear that κ_S is just a scaling factor, and the elasticity of i_S with respect to any policy change is thus independent of κ_S .

Second, the calibrated level of i_S also does not matter for policy. In all equilibrium conditions, outcomes do not depend on i_S , but on the product $x \cdot i_S$. If we were to abstract from the distribution of incumbents over quality and technology gaps, this product would be exactly pinned down by two other targeted moments: the contribution of entrants to growth ($\frac{\ln(\lambda) \cdot x \cdot i_S}{g}$) and the growth rate itself (g). The target for the implementation rate of startups only pins down i_S and x , but this decomposition is irrelevant for growth outcomes.

Of course, with heterogeneous incumbents, the argument becomes somewhat more complex. However, as Figure C.3 shows, the irrelevance of the target for the implementation probability of startups still continues to apply quantitatively.

D Multiproduct Extension

This section describes the equilibrium and calibration of the multiproduct model in Section 5.2 of the main text.

Model solution As firms can produce multiple products, the incumbent's value function in principle has an additional state variable in this extended model: the number of products produced, denoted by k . However, it is easy to show that the value function is linear in k ,

so that

$$v(k, \boldsymbol{\omega}, \mathbf{n}) = \sum_{j=1}^k v(\omega_j, n_j),$$

where $\boldsymbol{\omega} \equiv \{\omega_j\}_{j=1}^k$ and $\mathbf{n} \equiv \{n_j\}_{j=1}^k$ are the vectors containing the product quality and technology gap of every product produced by the firm.

Using this result, we can derive the HJB equation for the incumbent's value function in the extended model. To do so, we first derive the continuation values of incumbents and startups after the arrival of a startup idea. Consider the case of a startup that has an idea on a product j , characterized by a quality ω and a technology gap n . Proceeding by backward induction, suppose first that this startup does not meet the current producer of product j , and is matched to the current producer of product j' , characterized by a quality ω' and a technology gap n' . For this startup, the continuation value in case of no meeting with the unrelated incumbent is still given by equation (A.2):

$$v_S^{\text{NoMeet}}(\omega) = \max_{i_S} \left\{ i_S \cdot \left(\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot v(\omega, n_S) \right) - \kappa_S \cdot i_S^\psi \right\}. \quad (\text{D.1})$$

For the unrelated incumbent, instead, the continuation value in case of no meeting with the startup is just $v(\omega', n')$. The surplus generated by a meeting between the startup and the unrelated incumbent is therefore

$$\begin{aligned} \sigma_U(\omega) &= \max \left[0, v(\omega', n') + \max_{i_U} \left\{ i_U \cdot \left(\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot v(\omega, n_S) \right) - \kappa_I \cdot i_U^\psi \right\} - v(\omega', n') - v_S^{\text{NoMeet}}(\omega) \right] \\ &= \max \left[0, \max_{i_U} \left\{ i_U \cdot \left(\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot v(\omega, n_S) \right) - \kappa_I \cdot i_U^\psi \right\} - v_S^{\text{NoMeet}}(\omega) \right]. \end{aligned} \quad (\text{D.2})$$

Comparing the surplus with (D.1), note that the surplus $\sigma_U(\omega)$ is strictly positive if and only if $\kappa_I < \kappa_S$. There is no rent-preserving motive for acquiring unrelated startups: these acquisitions can only be motivated by lower implementation costs for the incumbent. Moreover, as the unrelated incumbent does not currently produce the product of the acquired startup, the Arrow replacement effect does not apply: acquisitions of unrelated startups always increase the probability that the startup's idea is implemented.

Note that the surplus only depends on startup product quality ω , and is independent of the characteristics of the acquirer. Thus, all acquirers choose the same implementation probability $i_U(\omega)$, and all incumbents choose the same search effort s_U for unrelated startups. The continuation value for an incumbent after meeting an unrelated startup with

an idea on a product with quality ω is

$$v^{\text{MeetU}}(\omega', n') = v(\omega', n') + \alpha \cdot \sigma_U(\omega). \quad (\text{D.3})$$

For a startup, in turn, the continuation value when meeting an unrelated incumbent is

$$v_S^{\text{MeetU}}(\omega) = v_S^{\text{NoMeet}}(\omega) + (1 - \alpha) \cdot \sigma_U(\omega). \quad (\text{D.4})$$

With this, we can consider the interaction between the startup and the related incumbent. In the absence of a meeting between these two firms, the value of the incumbent is given by

$$v^{\text{NoMeetR}}(\omega, n) = \left[1 - \left((1 - s_U) \cdot i_S(\omega) + s_U \cdot i_U(\omega) \right) \right] \cdot v(\omega, n). \quad (\text{D.5})$$

That is, the related incumbent is displaced with probability $i_S(\omega)$ if the startup is not acquired, and with probability $i_U(\omega)$ if the startup is acquired.

For the startup, in turn, the expected continuation value in case of no meeting with the related incumbent is

$$v_S^{\text{NoMeetR}}(\omega) = v_S^{\text{NoMeet}}(\omega) + s_U \cdot (1 - \alpha) \cdot \sigma_U(\omega). \quad (\text{D.6})$$

This is the sum of the value of the startup as a stand-alone firm, and the expected surplus from a potential meeting with an unrelated incumbent.⁴⁰

Finally, the surplus from a meeting between a startup and a related incumbent is

$$\sigma_R(\omega, n) = \max \left[0, v(\omega, n) + \max_{i_A} \left\{ i_A \cdot \left(\sum_{n_S=1}^{+\infty} \theta(n_S) \cdot v(\omega, n + n_S) - v(\omega, n) \right) - \kappa_I \cdot i_A^\psi \right\} - v^{\text{NoMeetR}}(\omega, n) - v_S^{\text{NoMeetR}}(\omega) \right]. \quad (\text{D.7})$$

As in the baseline model, the continuation value of the startup and the related incumbent in case of a meeting is then just the sum of their outside options and their share of the surplus.

With these elements, we can now write the HJB equation of an incumbent firm as

$$\rho \cdot v(\omega, n) = \max_{z, s_R, s_U} \left\{ \omega \cdot (1 - \lambda^{-n}) - \zeta_I \cdot z^\psi - \chi_R \cdot s_R^\varphi - \chi_U \cdot s_U^\varphi \right.$$

⁴⁰Note that, as in the baseline model, we are implicitly assuming here that all firms are risk-neutral.

$$\begin{aligned}
& + z \cdot \max_{i_I} \left[i_I \cdot \left(v(\omega, n+1) - v(\omega, n) \right) - \kappa_I \cdot i_I^\psi \right] \\
& + x \cdot \left[s_R \cdot \alpha \cdot \sigma_R(\omega, n) - \left((1 - s_U) \cdot i_S(\omega) + s_U \cdot i_U(\omega) \right) \cdot v(\omega, n) \right] \\
& + x \cdot s_U \cdot \mathbb{E}_{\omega', n'} \left((1 - s_R(\omega', n')) \cdot \alpha \cdot \sigma_U(\omega') \right) \Big\} \\
& + \sum_{\omega'} \tau_{\omega, \omega'} \cdot \left[v(\omega', n) - v(\omega, n) \right]. \tag{D.8}
\end{aligned}$$

Finally, in equilibrium, the value of a startup idea must hold

$$\tilde{\zeta}_S = \mathbb{E}_{\omega, n} \left(v_S^{\text{NoMeet}}(\omega) + (1 - \alpha) \cdot \left(s_U \cdot \sigma_U(\omega) + s_R(\omega, n) \cdot \sigma_R(\omega, n) \right) \right).$$

These conditions summarize the dynamic problem of incumbents and startups. In the next paragraph, we describe the remaining equilibrium conditions and the solution algorithm.

Closing the model To close the model, we need to derive the generator matrix for the invariant distribution of (ω, n) . This generator matrix has the same structure as the one in the baseline model, shown in Section A.2. The only difference is that now, unrelated incumbents might also implement startup ideas. For the technology gap, these operations are exactly equivalent to startups implementing their own ideas. Thus, to obtain the new distribution, it is sufficient to replace the probability with which startups implement their own idea $((1 - s(\omega, n)) \cdot i_S(\omega)$ in the baseline model) with the probability that either the startup or an unrelated incumbent implements the startup idea, given by

$$(1 - s_R(\omega, n)) \cdot (s_U \cdot i_U(\omega) + (1 - s_U) \cdot i_S(\omega)). \tag{D.9}$$

Likewise, the growth rate is still given by equation (17), where the probability that a startup implements its own idea is replaced with the probability that a startup or an unrelated incumbent implements the startup idea.

To solve this model, we cannot use the exact same algorithm as in the baseline model. This is because the solution of the VFI problem now depends not only on the startup rate x , but also on the entire distribution of products over quality and technology gaps (ω, n) . Thus, instead of our baseline loop over the startup rate, we now need a joint loop over the

startup rate and the invariant distribution.⁴¹

Calibration and model fit We calibrate the multiproduct model by targeting the same moments as in the baseline model, as well as the percentage of acquisitions in which the startup is unrelated to the incumbent. In our sample, 59% of acquired startups do not belong to the same SIC 3-digit industry than the acquiring incumbent, and we use this number as a proxy for the percentage of unrelated acquisitions.

Table D.1: Fit of the multiproduct model.

Targeted moment	Model	Data	Data source	Identifies
Growth rate	2.0%	2.0%	Jones (2016)	λ
Exit rate	7.3%	7.3%	BDS	ξ_S
Growth contribution of entrants	25.7%	25.7%	Akcigit and Kerr (2018)	ξ_I
Avg. implementation prob., startups	10.0%	10.0%		κ_S
Effect of acq. on implementation prob.	0.0374	0.0374	Section 3	κ_I/κ_S
Percentage of startups acquired	4.0%	4.0%	Section 3	χ_R
Relative size of acquiring firms	2.10	2.10	Section 3	φ
Percentage of unrelated acquisitions	59%	59%	Section 3	χ_U

Table D.1 summarizes the fit of the multiproduct model. Just as the baseline model, the model exactly matches all moments.

⁴¹Moreover, in order to compute the search effort for related startups, we need to know the search effort for unrelated startups, and vice-versa. Thus, we need to solve for these two decisions with an inner loop.