The Effects of Startup Acquisitions on Innovation and Economic Growth*

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

May 7, 2024

Abstract

Innovative startups are frequently acquired by large incumbent firms. Such acquisitions have recently come under scrutiny, as policymakers suspect that incumbents might acquire startups just to "kill" their ideas. However, acquisitions also provide an incentive for startup creation, and have ambiguous effects on incumbents' own innovation. Our paper assesses the net effect of these forces. To do so, we build an endogenous growth model with heterogeneous multi-product firms and startup acquisitions, and calibrate its parameters to match micro-level evidence from the United States. Our calibrated model implies that a startup acquisition ban lowers the startup rate, but increases incumbent innovation as well as the implementation rate of startup ideas. As the negative forces are slightly stronger, the ban lowers growth by 0.03 percentage points per year. These results crucially depend on transaction prices: in the presence of higher acquisition premia, bans would increase growth.

Keywords: Acquisitions, Innovation, Productivity growth, Firm dynamics.
JEL Classification: O30, O41, E22.

*Paolo Mengano, Sutanuka Roy and Raquel Sánchez provided excellent research assistance. We thank Igor Letina, Huiyu Li, Johannes Matt, and Ludovic Panon for helpful discussions. We also thank Gadi Barlevy, Vasco Carvalho, Laurent Cavenaile, Murat Celik, Joel David, Maarten De Ridder, Florian Ederer, Chiara Fumagalli, Jorge Guzman, Hugo Hopenhayn, Chad Jones, Danial Lashkari, Enrico Moretti, Jane Olmstead-Rumsey, Michael Peters, Joshua Weiss, Liangjie Wu, and seminar participants at Banco de España, Bank of Lithuania, Bocconi, CEMFI, Chicago Fed, Danmarks Nationalbank, EIEF, European Commission, FRB of St Louis, IMF, Jinan, Kent, PSE, Sabanci, UAB, UB, UC Berkeley, UC Santa Cruz, UPF-CREi, University of Toronto, SED 2021, EEA-ESEM 2021, EWMES 2021, the CEPR 4th Joint Research Conference, T2M 2022, the 2022 Bristol and Lisbon Macro Workshops, BSE Summer Forum 2022, MMF 2022, the V Workshop of the Spanish Macro Network, the Warwick-CFM-Vienna Global Macro Workshop, 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.

‡Banco de España and CEMFI; 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 and make a substantial contribution to aggregate productivity growth in the United States (see e.g. Decker *et al.*, 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, a bipartisan group of Representatives proposed the Platform Competition and Opportunity Act, aiming to prohibit acquisitions by 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 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 *et al.*, 2021), acquiring startups only to kill innovative ideas that threaten the incumbent's business. Industry lobbyists have pushed back and argued that acquisitions actually foster innovation, as 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, policymakers 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 aforementioned forces discussed in the public debate, and also highlights the less commonly discussed (but no less important) feedback effects of acquisitions on incumbents' own innovation incentives.

[1] For instance, according to the Federal Trade Commission (FTC), the Tech giants Amazon, Apple, Facebook, Google and Microsoft acquired more than 600 small firms between 2010 and 2019 (FTC, 2021).

[2] A 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.

[3] See https://itif.org/publications/2022/01/31/platform-competition-and-opportunity-act-solution-search-problem. The act was, however, never voted upon (see https://www.govtrack.us/congress/bills/117/hr3826).

The model builds on the Schumpeterian growth framework with heterogeneous firms (Klette and Kortum, 2004; Peters, 2020). There is a continuum of incumbents, each producing a finite number of differentiated products. Incumbents invest in innovation to increase their productivity for existing products as well as to overtake products from others. Moreover, a large mass of non-producing startups invest in innovation to displace incumbents and enter the market. We introduce two new elements into this setting. First, by investing into a search technology, incumbents can acquire startups. These acquisitions can either be "related" (an incumbent buys a startup with an idea on one of the incumbent's products) or "unrelated" (an incumbent buys a startup with an idea on a product that the incumbent does not produce). Second, we assume that startup ideas need additional implementation investments to be realized. This allows us to capture the fact that incumbents might be more or less likely to implement an idea than the startup that came up with it.

To regulate startup acquisitions, we assume that a government can impose a tax on them. In our model, the effect of changes in this tax on the growth rate can be decomposed into three margins: (i) changes in the startup rate (i.e., the number of startups created at any given time); (ii) changes in the own innovation rate of incumbents; and (iii) changes in the implementation rate of startup ideas.

A higher tax on startup acquisitions has a direct effect on the third margin, the implementation rate of startup ideas. However, the direction of this effect is unclear. On the one hand, incumbents have less incentives to implement a startup's idea if it displaces some of their pre-existing profits. This replacement effect (Arrow, 1962) explains why incumbents might kill some startup ideas. On the other hand, when incumbents have lower implementation costs, they might actually implement startup ideas at higher rates than the startups themselves.[4] Depending on the relative strength of these forces, limiting startup acquisitions could either increase or decrease the implementation rate of startup ideas.

However, a full assessment of this policy needs to go beyond these direct effects, and take into account general equilibrium spillovers on the startup rate and on incumbents' own innovation (the other two margins in our decomposition). In our model, startups only agree to an acquisition if the price that they receive exceeds their outside option of remaining independent.[5] Thus, limiting their opportunities to be acquired lowers the payoff to startup creation and hence the startup rate. Finally, the response of incumbents' own innovation

[4] In our model, implementation costs are exogenous. Cost differences could be micro-founded by economies of scale and scope, a larger customer base, a better access to capital markets, or greater business experience.

[5] In the business world, there is a widespread conviction that acquisitions are a desirable outcome for startups. Numerous guides advise entrepreneurs how to position their startups in order to be acquired (see, e.g., https://www.forbes.com/sites/alejandrocremades/2019/08/02/how-to-get-your-business-acquired.

is ambiguous. Acquisitions generate a surplus for incumbents, allowing them to avoid displacement and to acquire ideas. Thus, all else equal, banning them lowers incumbent value and innovation incentives. However, the lower startup rate induced by a ban also reduces creative destruction, which increases incumbent value and innovation incentives.

Our objective is to assess the quantitative importance of these various forces. To do so, we discipline the model by calibrating it to micro data on acquisitions and patenting.

We build a new dataset combining three sources: acquisition data from SDC Platinum, patent data from PatentsView, and accounting data on publicly listed firms from Compustat. We focus on the period 2000-2020, and define startup acquisitions as deals in which the target is younger than six years, following Guzman and Stern (2020). With this data, we study the causal effect of acquisitions on the implementation of startup ideas, and then use our empirical estimate as a calibration target for our model. To measure implementation, we track patent citations. We interpret an increase in the post-acquisition citations received by a startup patent 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 match each acquired patent to a group of control patents with similar characteristics. We find that the average acquisition has virtually no effect on citations. Through the lens of our model, this implies that incumbents must have lower implementation costs, compensating for the Arrow replacement effect. However, we do find a negative effect on citations if the startup and its acquirer belong to the same NAICS 3-digit industry. This is consistent with our model, which suggests that related acquisitions have a more negative effect on implementation than unrelated acquisitions.

We calibrate the model to match our regression evidence and other relevant data moments, such as the average acquisition premium or the frequency of startup acquisitions. Then, we study the effects of a ban on startup acquisitions (implemented through an arbitrarily high tax on them). We start with a partial equilibrium analysis that only takes into account the direct effect of the ban on the implementation rate of startup ideas. We find that this effect is positive: although the average acquisition has no effect on implementation (in line with our empirical results), related acquisitions, which lower the implementation probability, are concentrated among larger firms. As larger firms have a higher weight in aggregate outcomes, the positive effect of the acquisition ban on implementation dominates.

Once we take into account general equilibrium feedback effects, however, the positive effect of the acquisition ban is overturned. The ban lowers the startup rate by a substantial 16.8%. This is only partially compensated by a small increase in incumbent's own innovation and in the implementation efforts of startups (both driven by the fact that a lower creative destruction threat from startups increases incumbent value). Overall, the

growth rate falls by 0.03 percentage points (from 2% to 1.97% per year). These results are mainly driven by related acquisitions, which generate a larger surplus than unrelated ones. Thus, a ban limited to related acquisitions has very similar effects to an overall ban.

In the final part of the paper, we show how these quantitative results are driven by our calibration targets. In different re-calibrations of the model, we find that our target for the acquisition premium (the percentage difference between the acquisition price and the startup's outside option) plays a key role. With a higher acquisition premium, acquisitions are less valuable for incumbents, and threatening startups are more costly. Therefore, acquisition bans trigger a larger increase in incumbent value and innovation. Precisely, when targeting an acquisition premium above 90% (instead of the baseline value of 49%), a ban increases growth. Thus, selective bans in industries with high premia could actually be beneficial. Likewise, bans are less detrimental if entrants contribute little to growth, or if the average effect of acquisition on implementation is strongly negative.

Related literature There is a growing empirical literature on the effect of mergers and acquisitions (M&As) on innovation. Cunningham *et al.* (2021) show that in the US pharmaceutical industry, acquirers are more likely to stop research projects of acquired firms when these overlap with their own drug portfolio. Such killer acquisitions are more frequent for acquirers with a dominant market position. Seru (2014) and Haucap *et al.* (2019) also provide evidence for a negative effect of M&As on firm R&D. Phillips and Zhdanov (2013) instead argue that acquisitions stimulate innovation by small 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 (2022) and Liu (2022) study the effect of acquisitions 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 *et al.*, 2017; Cabral, 2018; Bourreau *et al.*, 2018; Bryan and Hovenkamp, 2020; Fumagalli *et al.*, 2020; Kamepalli *et al.*, 2020; Letina *et al.*, 2020; Denicolò and Polo, 2021; Brutti and Rojas, 2022; Callander and Matouschek, 2022; 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.[6] Lentz and Mortensen (2016) and Akcigit *et al.* (2016) incorporate a market for ideas (through buyouts or patent sales) in endogenous growth models, showing that this improves the allocation of ideas. Finally, Cavenaile *et al.* (2021) and Olmstead-Rumsey *et al.* (2024) use endogenous growth models to study mergers between incumbents or the innovation effects of cross-industry acquisitions by Tech platforms. Our focus is different: we aim to quantify the overall growth effect of startup acquisitions, taking into account new general equilibrium spillovers on the startup rate and on the implementation of ideas.[7] 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 study startup acquisitions.

Outline In the remainder of the paper, 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 the implementation of startup ideas. Section 4 presents our calibration strategy, and Section 5 discusses our quantitative results. Section 6 concludes.

2 Model

In this section, we develop a model of the links between startup acquisitions and innovation. We build on Schumpeterian growth models with heterogeneous firms, but introduce two new elements: startup acquisitions and a distinction between the invention and the implementation of startup ideas.

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

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

[6] 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).

[7] Weiss (2022) studies changes in innovation costs in an endogenous growth model, 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. Our paper is also related to Celik *et al.* (2022), who study the effects of information frictions in the merger market on innovation and business dynamism.

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 inelastically supplies $L > 0$ units of labor at the wage w_t. She also owns all the firms in the economy and accumulates wealth A_t according to the budget constraint $\dot{A}_t = r_t A_t + w_t L + T_t - C_t$, where r_t is the equilibrium rate of return on assets and T_t is a lump-sum transfer.

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

$$Y_t = \exp\left(\int_0^1 \omega_j \ln\left(y_{jt}\right) \mathrm{d}j\right), \tag{2}$$

where y_{jt} is the output of product j at instant t, and ω_j is the spending share of product j. Spending shares are positive parameters. They take values in a finite set Ω and hold $\sum_{\omega\in\Omega} \omega\tau(\omega) = 1$, where $\tau(\omega) \in [0,1]$ denotes the (exogenous) mass of products with spending share ω.

Each product can potentially be produced by a large number of multi-product firms. Each firm f can produce a product j with a linear technology that uses labor,

$$y_{jft} = q_{jft}\ell_{jft}, \tag{3}$$

where y_{jft} is the output of firm f at instant t, q_{jft} is the productivity of the firm, and ℓ_{jft} is the labor input. At every instant, firms engage in Bertrand competition on product markets. Therefore, in equilibrium, each product is only produced by the firm that has the highest productivity for it. We denote this highest productivity by q_{jt}.

Productivity is improved through innovations by incumbents (firms which already produce one or more products) and startups (firms that do not produce yet). The next section describes their innovation technologies.

Incumbent innovation Incumbent firms do research to improve their existing products and to overtake the production of new ones. For each product that it produces, an incumbent can generate a Poisson arrival rate z of internal innovations by paying a cost of $\zeta_I^{\text{int}} z^\psi Y_t$ units of the final good. In this cost function, $\zeta_I^{\text{int}} > 0$ is a scaling factor and ψ is the elasticity of innovation costs with respect to innovation. We assume $\psi > 1$, so that costs are increasing and convex in the arrival rate of innovation. Costs are also proportional to aggregate output Y_t, to ensure balanced growth. An internal innovation enables the incumbent to improve its product's productivity by a factor $\lambda > 1$.

An incumbent can also generate external innovations, which improve the productivity

of products currently produced by others. By investing $\xi_I^{\text{ext}} n x^{\psi} Y_t$ units of the final good, an incumbent that currently produces $n = 1, 2, \ldots$ products generates a Poisson arrival rate nx of external innovations. External innovations apply to a randomly chosen product in the interval $[0, 1]$, and enable the incumbent to improve its leading productivity by a factor λ. The incumbent then becomes the new producer of this improved product, displacing the previous producer through creative destruction.

Startup innovation Innovation is also generated by startups. At every instant, there is a large mass of startups, which can pay a cost of $\xi_S Y_t$ units of the final good to generate a Poisson arrival rate 1 of ideas. A startup's idea applies to a randomly chosen product $j \in [0, 1]$. We assume that startup ideas are rough and need costly implementation. When the startup invests $\kappa_S i_S^{\theta} Y_t$ units of the final good (with $\kappa_S > 0$ and $\theta > 1$), it implements the idea with probability i_S.[8] An idea that is not implemented disappears forever. An implemented idea increases the leading productivity of the product by a factor λ, allowing the startup to leapfrog the current producer, enter and become an incumbent.

Startup acquisitions The startup may not always choose to invest into implementation: alternatively, it can be acquired by an incumbent. Acquisitions can take place if, and only if, there is a meeting between the startup and the potential acquirer. For such a meeting to happen, two conditions must be met: first, the startup needs to be matched to the acquirer; and second, the acquirer needs to notice the startup.

As described before, the startup's idea applies to a product $j \in [0, 1]$. We assume that with probability $1 - \gamma$, the startup is matched to the related incumbent, i.e., to the current producer of product j. With probability γ, however, it is matched to an unrelated incumbent, i.e., to the producer of another product j', randomly chosen from the interval $[0, 1]$.[9]

The probability that potential acquirers notice the startup is endogenous, and depends on their effort in monitoring the startup scene.[10] We assume that for each of its products, an incumbent spends $\chi s^{\varphi} Y_t$ (with $\chi > 0$ and $\varphi > 1$) units of the final good to generate a probability s of noticing a matched startup. This needs to be done separately for related and unrelated startups: that is, the incumbent chooses a probability s_R for noticing related startups and a probability s_U for noticing unrelated startups. Just like innovation costs,

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

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

[10] We think of this as a reduced-form model of information and search frictions in the acquisition market, preventing incumbents from noticing all startups and forcing them to spend resources to monitor the market.

search costs are increasing and convex, and scale with aggregate output to ensure balanced growth.

When there is a meeting, the incumbent may acquire the startup. To do so, it needs to pay a flat acquisition tax $T_R Y_t$ (for a related acquisition) or $T_U Y_t$ (for an unrelated acquisition).[11] The acquisition occurs if it generates a surplus, that is, if the joint value of both firms after the acquisition is greater than the sum of their outside options and the acquisition tax. 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: by investing $\kappa_I i^\theta Y_t$ units of the final good, where $\kappa_I > 0$, it implements the startup's idea with probability i. Successful implementation allows the incumbent to improve its existing product (if the startup is related) or to add a new product to its portfolio (if the startup is unrelated).[12]

2.2 Equilibrium

2.2.1 Household decisions, markups 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 $g > 0$. On the BGP, consumption growth holds the Euler equation

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

The demand function for every product j is given by $y_{jt} = \omega_j \frac{Y_t}{p_{jt}}$. Bertrand competition implies that for any product j, only the firm with the highest productivity produces in equilibrium. However, its pricing decisions depend on the firm with the second-highest productivity. This "follower" is an old incumbent, displaced when the current incumbent overtook production. Denoting by q_{jt}^F the follower's productivity, we define the technology gap (the number of productivity steps between the incumbent and the follower) as the integer a_{jt} holding $\lambda^{a_{jt}} \equiv q_{jt}/q_{jt}^F$. Limit pricing implies that the incumbent charges a price equal to the marginal cost of the follower, in order to keep that firm out of the market.[13] As

[11] We assume that the proceeds from these acquisitions are rebated to the representative consumer.

[12] Figures A.1 and A.2 in Appendix A.4 summarize the timing of events for startup ideas.

[13] In general, the optimal price is the minimum between the marginal cost of the follower and the unconstrained monopoly price. However, because the price elasticity is equal to one, the latter is infinite.

the follower's marginal cost exceeds the incumbent's by a factor $\lambda^{a_{jt}}$, we have

$$p_{jt} = \mu(a_{jt})\frac{w_t}{q_{jt}}, \text{ with } \mu(a_{jt}) = \lambda^{a_{jt}}, \tag{5}$$

where $\mu(a_{jt})$ denotes the markup. These results imply that an incumbent producing a product with spending share ω_j and technology gap a_{jt} makes a profit

$$\pi_t\left(\omega_j, a_{jt}\right) = \omega_j\left(1-\lambda^{-a_{jt}}\right)Y_t \tag{6}$$

on that product. Profits are increasing linearly in the spending share, and increasing and concave in the technology gap. They do not depend on the productivity level q_{jt}.

2.2.2 Dynamic problem of incumbent firms

For each of their products, incumbents choose an internal innovation rate z, as well as search efforts s_R (for related startups) and s_U (for unrelated startups). Moreover, they decide whether to acquire a startup whenever they meet one, an implementation probability for the startup's idea, as well as an external innovation rate x.

The value function of an incumbent depends on the number of products it produces, and on the spending share and technology gap of each product. However, the fact that external innovations scale linearly in firm size (as in Klette and Kortum, 2004) implies that the value function is additively separable across products. Moreover, as profits and all costs are proportional to aggregate output, the value function is proportional to aggregate output, too. Thus, as we prove in Appendix A.1, the value function $V_t(\mathbf{n}, n)$ of a firm with n products and a portfolio $\mathbf{n} \equiv \left(\omega_j, a_j\right)_{j=1}^{n}$ of spending shares and technology gaps, holds

$$V_t(\mathbf{n}, n) = \sum_{j=1}^{n} v(\omega_j, a_j)Y_t \tag{7}$$

for some time-invariant value function v. The Hamilton-Jacobi-Bellman (HJB) equation for the value function of an individual product in state (ω, a) holds

$$\rho v(\omega,a) = \max_{\substack{z,x \\ s_R,s_U}} \Bigg\{ \underbrace{\omega(1-\lambda^{-a})}_{\text{Profits}} - \underbrace{\zeta_I^{\text{int}} z^{\psi}}_{\substack{\text{Internal} \\ \text{inn. cost}}} - \underbrace{\zeta_I^{\text{ext}} x^{\psi}}_{\substack{\text{External} \\ \text{inn. cost}}} - \underbrace{\chi\left(s_R^{\varphi} + s_U^{\varphi}\right)}_{\substack{\text{Search costs for rel.} \\ \text{and unrel. startups}}} + \underbrace{z\Big[v(\omega,a+1) - v(\omega,a)\Big]}_{\text{Own product innovation}}$$

$$+ \underbrace{x \sum_{\omega' \in \Omega} \tau(\omega') v(\omega',1)}_{\text{External innovation}} - \underbrace{\widetilde{x} v(\omega,a)}_{\substack{\text{Creative destruction} \\ \text{by external innovation}}} + \underbrace{x_S v_R(s_R,\omega,a)}_{\substack{\text{Startup has idea on} \\ \text{incumbent's product}}} + \underbrace{x_S \gamma v_U(s_U)}_{\substack{\text{Unrelated startup} \\ \text{matched to incumbent}}} \Bigg\}, \tag{8}$$

In this equation, x_S stands for the arrival rate of startup ideas and $\widetilde{x}$ for the arrival rate of other incumbents' external innovations. Both of these rates are aggregate endogenous outcomes, but they are taken as given by each individual incumbent.

The HJB equation shows how the discounted value of a product of type (ω, a) changes over time. First, at every instant, the incumbent collects profits and spends on innovation costs and startup search. Then, the incumbent makes an internal innovation at Poisson rate z, allowing it to increase its technology gap by one step. At rate x, the incumbent makes an external innovation, allowing it to overtake a new product. This new product will have a spending share ω', chosen at random from the distribution of spending shares, and a technology gap 1. At rate $\widetilde{x}$, the firm loses its product to another incumbent. Finally, the last two terms capture changes in value due to startups. With an arrival rate x_S, a startup has an idea on the incumbent's product. The startup might then try to implement this idea, be acquired by the incumbent, or be acquired by another, unrelated firm. This entails an expected change in the incumbent's value which we denote by $v_R(s_R, \omega, a)$, and will derive below. Finally, at rate $x_S\gamma$, the incumbent is matched to an unrelated startup. In that case, incumbent value increases by $v_U(s_U)$, which we will again derive below.

From problem (8), we find that the product-level internal innovation rate holds

$$z(\omega, a) = \left(\frac{v(\omega, a+1) - v(\omega, a)}{\zeta_I^{\text{int}} \psi} \right)^{\frac{1}{\psi - 1}}. \tag{9}$$

The external innovation choice, in turn, does not depend on current product characteristics:

$$x = \left(\frac{\sum_{\omega' \in \Omega} \tau(\omega') v(\omega', 1)}{\zeta_I^{\text{ext}} \psi} \right)^{\frac{1}{\psi - 1}}. \tag{10}$$

2.2.3 Unrelated acquisitions

Next, we derive the changes in incumbent value due to startups, v_U and v_R. These will allow us to characterize search and acquisition decisions.

First, consider the case in which an unrelated startup is matched to the incumbent. This startup has an idea on a product j', characterized by a spending share ω' and technology gap a'. When this match happens, there are two possibilities: either the incumbent notices the startup (with probability s_U), or it does not notice it. In the latter case, the change in the incumbent's value is zero. In the former case, however, the incumbent might acquire the startup and implement its idea, thereby increasing its value.

To determine whether this happens, we compute the surplus from an unrelated acquisi-

tion. For this, we first need to compute the incumbent's and the startup's outside options. For both firms, the outside option is the change in their continuation value in the absence of a meeting. For the incumbent, this is zero. For the startup, the outside option is the value of the problem

$$\max_{i_S} \left\{ i_S v(\omega', 1) - \kappa_S i_S^{\theta} \right\}. \tag{11}$$

Without a meeting, a startup chooses an optimal implementation probability i_S. When its idea is implemented, the startup enters and becomes the new incumbent producer of product j'. Instead, if implementation fails, the startup has a continuation value of zero. This problem implies that the startup chooses

$$i_S(\omega') = \left(\frac{v(\omega', 1)}{\kappa_S \theta} \right)^{\frac{1}{\theta - 1}}. \tag{12}$$

This probability does not depend on the technology gap a'. Plugging it back into equation (11), the startup's outside option becomes $\kappa_S (\theta - 1) \left(i_S(\omega') \right)^{\theta}$. The surplus from an unrelated acquisition, normalized by aggregate output, is then the difference between the incumbent's change in value when acquiring the startup's idea (net of the acquisition price), and the two outside options and the acquisition tax:[14]

$$b_U(\omega') = \underbrace{\max_{i_U} \left\{ i_U v(\omega', 1) - \kappa_I i_U^{\theta} \right\}}_{\text{Incumbent value change after acquisition}} - T_U - \underbrace{0}_{\substack{\text{Incumbent} \\ \text{outside option}}} - \underbrace{\kappa_S (\theta - 1) \left(i_S(\omega') \right)^{\theta}}_{\substack{\text{Startup} \\ \text{outside option}}}, \tag{13}$$

The incumbent implements the acquired idea with probability

$$i_U(\omega') = \left(\frac{v(\omega', 1)}{\kappa_I \theta} \right)^{\frac{1}{\theta - 1}}. \tag{14}$$

Just as in the case of the startup's policy, this implementation probability only depends on the product's spending share ω'. Using this, the surplus in (13) can be simplified to

$$b_U(\omega') = (\theta - 1) \left(\kappa_I \left(i_U(\omega') \right)^{\theta} - \kappa_S \left(i_S(\omega') \right)^{\theta} \right) - T_U. \tag{15}$$

An acquisition on an unrelated startup occurs if, and only if, this surplus is positive. Without taxes, this is the case when the incumbent has lower implementation costs than the startup (i.e., $\kappa_I < \kappa_S$). In this case, the idea is more valuable in the hands of the incumbent, and an

[14] The acquisition price is a transfer between the two firms. Thus, it does not matter for the surplus.

acquisition allows this transfer to take place. We can finally write the expected change in the incumbent's value after being matched to an unrelated startup as

$$v_U(s_U) = s_U \alpha \sum_{\omega' \in \Omega} \tau(\omega') \max\Big(0, b_U(\omega')\Big). \tag{16}$$

In words, when the incumbent notices the startup (with probability s_U), the change in value is the incumbent share α of the acquisition surplus. Taking the expectation over all products to which an unrelated startup's idea may apply, we obtain the expression above.

From this, it finally follows that the search effort for unrelated startups is

$$s_U = \left[\frac{\alpha \gamma x_S}{\chi \varphi} \left(\sum_{\omega' \in \Omega} \tau(\omega') \max\left(0, b_U(\omega')\right) \right) \right]^{\frac{1}{\varphi - 1}}. \tag{17}$$

The search effort for unrelated startups is independent of the incumbent's current characteristics. It is increasing in the expected acquisition surplus, the arrival rate of startup ideas, and the incumbent's bargaining weight.

2.2.4 Related acquisitions

Next, we derive the incumbent's change in value when a startup has an idea on its product. When this related incumbent is not matched to the startup, it faces potential displacement either from the startup itself (with probability $1 - s_U$) or from an unrelated incumbent acquiring the startup (with probability s_U). The startup implements with probability $i_S(\omega)$, and the unrelated incumbent with probability $i_U(\omega)$.[15] In both cases, implementation makes the related incumbent lose its entire value. Hence, its value decreases in expectation by $[s_U i_U(\omega) + (1 - s_U) i_S(\omega)] v(\omega, a)$.

In turn, when the related incumbent is matched to the startup, it can notice the startup (with probability s_R) or not notice it (with probability $1 - s_R$). In the latter case, the expected value of the incumbent decreases by $i_S(\omega) v(\omega, a)$, as the incumbent is displaced with probability $i_S(\omega)$. Alternatively, when the incumbent notices the related startup, it might acquire it. To characterize its decision, we need to again compute the acquisition surplus. The outside option of the startup is still given by the solution to problem (11), for a product with spending share ω. The outside option of the incumbent, in turn, is its change in value in case it did not notice the startup. Given this, the expected surplus from

[15] Note that s_U and $i_U(\omega)$ are taken as given by the related incumbent. However, in equilibrium, they are equal to the incumbent's own choices. To simplify notation, we directly impose this equality here.

the acquisition of a related startup holds

$$b_R(\omega,a) = \underbrace{\max_{i_R}\left\{i_R\Big(v(\omega,a+1)-v(\omega,a)\Big)-\kappa_I i_R^{\theta}\right\}}_{\text{Incumbent change in value after a related acquisition}} - T_R \tag{18}$$
$$\ldots - \underbrace{\left(-i_S(\omega)\,v(\omega,a)\right)}_{\text{Incumbent outside option}} - \underbrace{\kappa_S\,(\theta-1)\,(i_S(\omega))^{\theta}}_{\text{Startup outside option}}.$$

The incumbent's change in value after the acquisition (expressed again net of the acquisition price) comes from potentially implementing the startup's idea and increasing its technology gap by one unit. The incumbent implements the idea of a related startup with probability

$$i_R(\omega,a) = \left(\frac{v(\omega,a+1)-v(\omega,a)}{\kappa_I\theta}\right)^{\frac{1}{\theta-1}}. \tag{19}$$

Plugging this back into equation (18), the surplus from a related startup acquisition is

$$b_R(\omega,a) = i_S(\omega)v(\omega,a) - T_R + (\theta-1)\left(\kappa_I\,(i_R(\omega,a))^{\theta} - \kappa_S\,(i_S(\omega))^{\theta}\right). \tag{20}$$

An acquisition takes place if, and only if, this surplus is positive. Then, the surplus is split between both firms according to their Nash bargaining weights.

Putting all of the above results together, we can finally write down the expected change in the value of the related incumbent when a startup has an idea on its product:

$$v_R(s_R,\omega,a) = -\Big(\gamma s_U i_U(\omega) + (1-\gamma s_U)\,i_S(\omega)\Big)v(\omega,a) + s_R\alpha\,(1-\gamma)\max\Big(0,b_R(\omega,a)\Big). \tag{21}$$

The first term is the change in value without an acquisition. The term inside the parentheses is the displacement probability, a weighted average of the implementation probability of startups and unrelated incumbents. The second term is the expected surplus from an acquisition, which avoids displacement. From this, the search effort for related startups holds

$$s_R(\omega,a) = \left(\frac{\alpha\,(1-\gamma)\,x_S\max(0,b_R(\omega,a))}{\chi\varphi}\right)^{\frac{1}{\varphi-1}}. \tag{22}$$

The search probability is increasing in the startup rate and in the acquisition surplus.

To conclude, replacing all our results into equation (8), product value holds

$$v(\omega,a) = \frac{\pi(\omega,a) + (\psi-1)\left(\xi_I^{\text{int}} z(\omega,a)^{\psi} + \xi_I^{\text{ext}} x^{\psi}\right) + \chi\,(\varphi-1)\left(s_R(\omega,a)^{\varphi} + s_U^{\varphi}\right)}{\rho + x + x_S\Big(\gamma s_U i_U(\omega) + (1-\gamma s_U) i_S(\omega))\Big)}, \tag{23}$$

where we have used the fact that in equilibrium, $\widetilde{x} = x$ (all incumbents choose the same external innovation rate). Firm value stems from profits and the values of internal innovation, external innovation and search for startups. It is discounted at the sum of the discount rate ρ and the displacement rate by startups and other incumbents.

Before proceeding, it is useful to compare the implementation probabilities of startups and incumbents. For related incumbents, we have

$$\left(\frac{i_R(\omega,a)}{i_S(\omega)}\right)^{\theta-1} = \underbrace{\frac{v(\omega,a+1)-v(\omega,a)}{v(\omega,1)}}_{\text{Arrow replacement effect}} \quad \underbrace{\frac{\kappa_S}{\kappa_I}}_{\substack{\text{Cost}\\ \text{differences}}} \tag{24}$$

As we will show later, the value function is concave in the technology gap. Thus, the first term is always smaller than 1, reflecting the Arrow replacement effect: innovation implies a smaller increase in value for an incumbent (who displaces some preexisting profits) than for a startup.[16] All else equal, this implies that related acquisitions lower the implementation probability of startup ideas. However, there is also a cost effect: if the incumbent has lower implementation costs, the acquisition may still increase the implementation probability.

On the other hand, the implementation probability for unrelated incumbents holds $\left(i_U(\omega)/i_S\right)^{\theta-1} = \kappa_S/\kappa_I$. Startups and incumbents do not have pre-existing profits for unrelated products, so that they only make different choices if they have different costs.

2.2.5 Startup creation

Next, we determine the equilibrium startup rate, x_S. In an equilibrium with startup creation, the cost of creating a startup must be equal to the expected benefit of creating it:

$$\xi_S = \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m(\omega,a)\Bigg[\underbrace{\kappa_S\,(\theta-1)\,(i_S(\omega))^{\theta}}_{\text{Startup's outside option}} + \dots$$

[16] The displacement effect does not apply only to profits. Each product represents "knowledge capital", as it increases innovation possibilities (Klette and Kortum, 2004). For the related incumbent, implementing the startup's idea is an internal innovation, which does not increase knowledge capital. However, for the startup, it is an external innovation, increasing its knowledge capital. Appendix A.1 contains further details on this.

$$\ldots + (1-\alpha) \underbrace{\left((1-\gamma)\, s_R(\omega,a) b_R(\omega,a) + \gamma s_U b_U(\omega) \right)}_{\text{Expected surplus from related and unrelated acquisitions}} \Bigg]. \quad (25)$$

To understand this equation, note that a startup's idea falls on a random product with spending share ω and technology gap a. The right-hand side of equation (25) takes an expectation using the joint distribution of products over these two states, denoted by $m(\omega, a)$. This distribution is an endogenous equilibrium object that we derive in Appendix A.2. The startup always obtains its outside option, and if it is noticed by an incumbent, it also obtains a share $(1-\alpha)$ of the acquisition surplus.[17] Thus, all else equal, acquisitions increase the value of creating a startup and therefore the startup rate.

2.2.6 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 $\frac{\omega_j}{\mu(a_j)} \frac{Y_t}{w_t}$ units of labor. Using labor market clearing, we obtain the aggregate labor share:

$$\frac{w_t L}{Y_t} = \frac{1}{\mathcal{M}}, \qquad \text{where } \mathcal{M} \equiv \left(\int_0^1 \omega_j \mu_{jt}^{-1} \mathrm{d}j \right)^{-1}, \quad (26)$$

and recall that $\mu_{jt} = \lambda^{a_{jt}}$. The aggregate labor share is the inverse of the aggregate markup, $\mathcal{M}$, defined as a harmonic weighted average of product-level markups.

From the definition of the aggregate price index, we obtain that the wage holds

$$w_t = \frac{Q_t}{\mathcal{W}}, \quad \text{where} \quad Q_t \equiv \exp\left(\int_0^1 \omega_j \ln(\omega_j q_{jt}) \mathrm{d}j \right). \quad (27)$$

In this equation, Q_t is an aggregate productivity index, and $\mathcal{W} \equiv \exp\left(\int_0^1 \omega_j \ln(\mu_{jt}) \mathrm{d}j \right)$. Product market clearing, in turn, implies that aggregate output is fully used for consumption (C_t), innovation by incumbents and startups (I_t), and search (S_t), so that:

$$Y_t = C_t + I_t + S_t. \quad (28)$$

In Appendix A.3, we show that I_t and S_t are both linear functions of aggregate output. Therefore, aggregate consumption C_t grows at the same rate as output. Moreover, we can write aggregate output as

$$Y_t = \mathcal{M} \mathcal{W}^{-1} Q_t L. \quad (29)$$

[17] We drop the max operator on the acquisition surplus, as a negative surplus implies a search effort of zero.

The term $\mathcal{MW}^{-1} < 1$ is an efficiency wedge, due to heterogeneous markups triggering a misallocation of labor.

Finally, we derive the growth rate of the economy. Equation (29) implies that $g = \frac{\dot{Y}_t}{Y_t} = \frac{\dot{Q}_t}{Q_t}$. Then, as we show in Appendix A.3,

$$g = \ln(\lambda) \underbrace{\left(x + \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} \left[m(\omega, a) \omega \left(z(\omega, a) + z_S(\omega, a) \right) \right] \right)}_{\text{Aggregate innovation rate}}, \quad \text{with} \tag{30}$$

$$z_S(\omega, a) \equiv x_S \Big((1-\gamma)\, s_R(\omega, a) i_R(\omega, a) + \gamma s_U i_U(\omega) + \big(1 - (1-\gamma)\, s_R(\omega, a) - \gamma s_U\big) i_S(\omega) \Big). \tag{31}$$

Growth is proportional to the aggregate innovation rate. The innovation rate is, in turn, an average of the own innovation rates of incumbents (x and z) and the innovation rate from startup ideas, z_S. Rates are weighted by spending shares ω (which in equilibrium are equal to product sales shares). Finally, the innovation rate from startup ideas is the product of the arrival rate of startup ideas and the probability that those ideas are implemented, either by incumbents that acquire them or by the startups themselves.

This completes the characterization of our model's BGP equilibrium. Appendix C explains in detail how we compute this equilibrium numerically.

2.3 Key properties of the model

Having characterized the BGP equilibrium, we now discuss some of its key properties.[18]

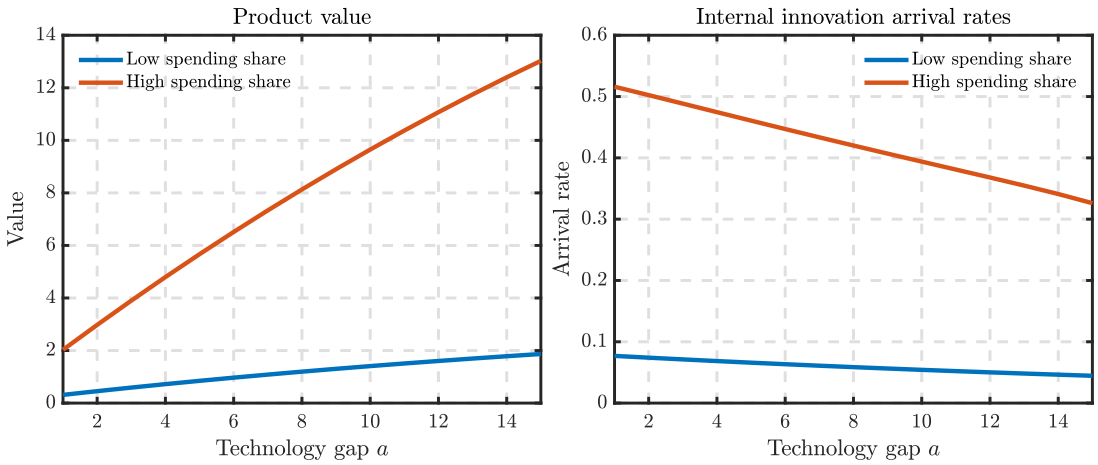

Figure 1: Value and internal innovation arrival rates by product type.

[18] All figures in this section have been drawn using our baseline calibration, discussed in Section 4.

Figure 1 plots the product-level value function v and the internal innovation arrival rate z for all products. We assume, as in our baseline calibration, that spending shares only take two values, a high and a low one. Product value is increasing in the spending share ω and in the technology gap a. Moreover, product value is concave in a, 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, it is increasing in the spending share ω and decreasing in the technology gap a.

The left panel of Figure 2 plots the noticing probabilities s_R and s_U as a function of product characteristics. For unrelated startups, the noticing probability is constant. For related startups, noticing probabilities are increasing in the spending share ω and the technology gap a. Indeed, related acquisitions have two benefits. First, when incumbents have lower implementation costs ($\kappa_I < \kappa_S$), they transfer an idea to a more efficient user. Second, they avoid business stealing and allow the technology gap a to remain at least at its current value, instead of being potentially lowered through entry. The importance of both effects is increasing in ω and a, explaining the behavior of noticing probabilities.

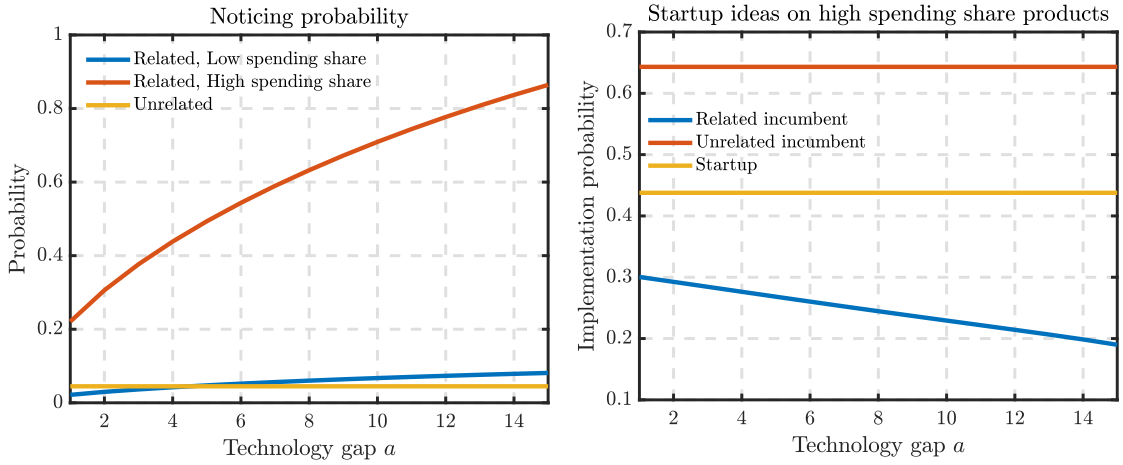

Figure 2: Noticing and implementation probabilities, by product type.

The right panel of Figure 2 plots the implementation probabilities for a startup idea on a high-spending-share product, distinguishing between the cases in which the startup invests into implementation itself, or is acquired by a related or unrelated incumbent. In our baseline calibration, incumbents have lower implementation costs than startups. Thus, unrelated incumbents are more likely to implement than the startup. In contrast, due to a strong replacement effect, related incumbents have lower implementation incentives in our calibration. Hence, the average related acquisition is a killer acquisition. Moreover, the marginal benefit of implementation for related incumbents decreases in the technology gap,

so that incumbents with high technology gaps are even more likely to do killer acquisitions.

The previous discussion shows that incumbent decisions about research, implementation and startup search crucially depend on the technology gap a. Therefore, the distribution of technology gaps across products, shown in Figure 3, is an important equilibrium object. This distribution is endogenous, shaped by the innovation and acquisition choices of incumbents and startups. In particular, the figure shows that high-spending-share products have on average higher technology gaps, as they receive more innovation.

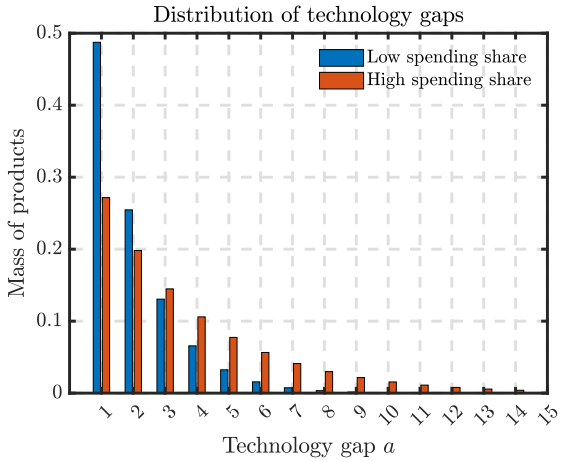

Figure 3: Technology gap distribution. *Notes:* This figure plots the distribution of technology gaps a across products, within each spending share category. Appendix A.2 describes how we solve for this distribution.

2.4 Decomposing the effects of startup acquisition policies on growth

In Section 5, we will use our model to study changes in antitrust policy, modeled as changes in startup acquisition taxes. These changes affect all endogenous variables in our model. To organize our discussion of these effects, we now show a useful decomposition result on the rate of economic growth, which we will repeatedly use in our analysis.

As we show in Appendix A.3, the change in the growth rate between the baseline BGP equilibrium A and an alternative BGP equilibrium B can be expressed as:

$$\begin{aligned}\frac{g^B}{g^A} = \text{Share}^A_{\text{inc}} \cdot \frac{\text{Incumbent own innovation}^B}{\text{Incumbent own innovation}^A} \\ + (1 - \text{Share}^A_{\text{inc}}) \cdot \left(\frac{\text{Startup rate}^B}{\text{Startup rate}^A} \cdot \frac{\text{Implementation rate of startup ideas}^B}{\text{Implementation rate of startup ideas}^A} \right),\end{aligned} \tag{32}$$

with

$$\text{Incumbent own innovation} = x + \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} \left[m(\omega, a) \omega z(\omega, a) \right],$$

$$\text{Startup rate} = x_S,$$

$$\text{Implementation rate of startup ideas} = \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} \left[m(\omega, a) \omega \frac{z_S(\omega, a)}{x_S} \right].$$

Equation (32) shows that in order to quantify the effect of any change in policy regarding startup acquisitions, we need to know the response of the three margins shown in the equation: (i) changes in incumbent own innovation (the sales-weighted average of incumbent internal and external innovation rates); (ii) changes in the startup rate; and (iii) changes in implementation rate of startup ideas (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 incumbent's own innovation, $\text{Share}^A_{\text{inc}}$.

All three margins shown in equation (32) are affected by startup acquisition taxes. In order to quantify these margins, we map the model to the data. To do so, the next section introduces a new micro-level data set, which allows us to generate calibration targets.

3 Empirical analysis

3.1 Data

We use three data sources: acquisition data from SDC Platinum, patent data from PatentsView, and accounting data for publicly listed firms from Compustat. This section describes the data and explains how we merge them. Appendix B.1 contains further details.

Data sources To track startup acquisitions, we rely on SDC Platinum, provided by the London Stock Exchange Group.[19] SDC provides transaction-level data on mergers and acquisitions. It includes information on the deal (e.g., the deal date and transaction value) and on the involved firms (e.g., their names, incorporation dates and industries). For our analysis, we consider transactions between 2000 and 2020 in which both the target and the acquirer firm are incorporated in the United States. Following Guzman and Stern (2020), we define startups as firms that are younger than 6 years. Thus, we further limit our sample to deals in which the target firm is younger than 6 years at the time of the acquisition.

In order to measure the innovation activity of firms, we rely on patent data from the USPTO's PatentsView database.[20] PatentsView contains the universe of US patents since 1976. It provides information on the firm owning each patent (name and assignee

[19] SDC is the standard data source for research on M&As (e.g. David, 2020; Guzman and Stern, 2020). It can be accessed at https://www.lseg.com/en/data-analytics/products/sdc-platinum-financial-securities.

[20] The data can be accessed at https://patentsview.org/.

code), as well as forward and backward citations, the number of claims assigned to the patent, the identify of the inventors, and patent technology classes.

Finally, we use Standard & Poor's Compustat North America database, which contains balance sheet and income statement information for all publicly listed firms in the U.S.[21]

Merging data sources Our baseline dataset is the sample of startup acquisitions from SDC Platinum. We use the other two data sources to add information to this baseline.

First, for each publicly listed acquirer, we use Compustat to infer some additional balance sheet data (e.g., acquirer sales in the acquisition year). For this, we rely on the crosswalk of Ewens *et al.* (2023), which links SDC deal numbers to firm identifiers in Compustat.

Second, we aim to match each acquired startup to its patents (if it holds any). This match is challenging, as there is no crosswalk between SDC Platinum and PatentsView, and the only common information is the name of the firm. To carry out the match, we rely on a series of routines developed by Wasi and Flaaen (2015) for Stata. We first standardize company names (using `stnd_compname`). Then, we use a fuzzy name matching algorithm (`reclink2`) to match firm names in SDC to firm names in the PatentsView. We complement this automatic match with a large scale manual check (see Appendix B.1 for details).

3.2 Descriptive statistics

Figure 4 plots the number of startup acquisitions per year in our final sample.

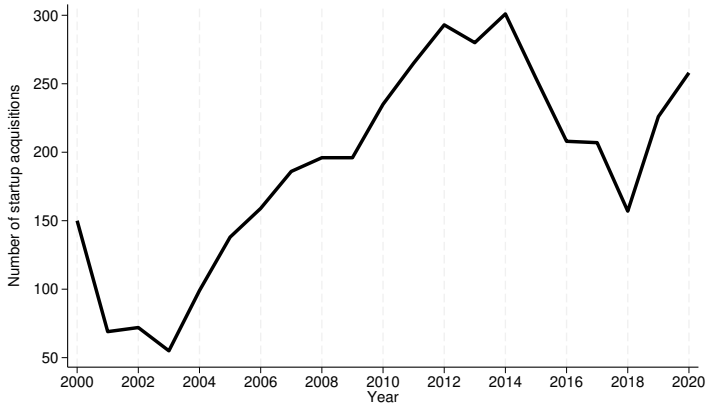

Figure 4: Number of startup acquisitions per year

After a low point in 2003 (with about 50 deals), acquisitions steadily increased until 2014, reaching almost 300 deals per year. After 2014, there has been a slowdown.[22]

[21] The data is at https://www.spglobal.com/marketintelligence/en/?product=compustat-research-insight.
[22] In an earlier version of this paper, we used data from Guzman and Stern (2020), who compiled a database

Table 1 summarizes some other key properties of our data set. The median startup is about 4 years old when it is acquired. SDC Platinum provides the deal value for about a quarter of all acquisitions. The median deal value is $31 million, but the distribution is highly skewed, with the mean exceeding the median by a factor of 10.[23] A subset of acquired startups (around 10% of those with a deal value) were already publicly listed before their acquisition. For these, SDC provides also the acquisition premium, that is, the percentage difference between the deal value and the startup's pre-acquisition stock market value (measured four weeks before the announcement of the deal). The average acquisition premium in our dataset is 49.3%.[24]

Table 1: Summary statistics for the startup acquisition sample

	N	Mean	Median	p10	p90
Acquisition year	4004	2012	2012	2004	2019
Target age (years)	4004	3.6	3.8	1.6	5.6
Deal value (millions)	1003	306	31	2	615
Acquisition Premium (percent)	104	49.3	28.8	-0.1	129.4
Same industry	4004	0.45	0	0	1
Acquirer sales (millions)	1183	13915.52	1159	36	37905
Patent Stock (Target Firms)	224	3.40	2	1	8
Stock Citations Received (Acquired Patents)	596	4.15	0	1	12
Stock Citations Received (Non-Acquired Patents)	1,445,115	1.21	0	0	2

Notes: This table considers a sample of deals between 2000 and 2020 in which both acquirer and target are incorporated in the US, and the target is younger than 6 years when the deal takes place. The variable "Same industry" is a dummy equal to 1 if target and acquirer share the same NAICS 3-digit code, and equal to 0 otherwise. The variable "Acquirer sales" is only available for acquirers that can be matched to Compustat. The variable "Patent Stock (Target Firms)" shows the number of patents for target firms that do have patents. The variable "Stock Citations Received" measures the number of citations received by target patents and non-target patents in the initial 5 years.

For our empirical analysis, we define related acquisitions as acquisitions in which the target and the acquirer share the same NAICS 3-digit industry code. As Table 1 shows, this is the case for 45% of startup acquisitions in our sample. Finally, we find that 224 target firms hold patents (with a median number of two). Measured by citations, these patents are of higher quality than the average patent: on average, they received 4.15 citations during their first five years, as opposed to 1.21 citations for all other patents.

on all new firms incorporated in 32 US states between 1988 and 2008. Their data indicates that about 4% of all patenting startups are acquired within their first six years. This percentage has no clear time trend.

[23] The deal with the highest value is Facebook's acquisition of WhatsApp for $19 billion.

[24] This number is similar to David (2020), who uses a broader sample of M&As between publicly listed firms and finds an average acquisition premium of 47%.

Acquirer characteristics Both the acquiring firms and the startups that they acquire are a selected sample of the overall population of firms. For instance, acquiring firms are generally larger than average. To show this formally, we consider the sample of all Compustat firms, and run a simple linear regression relating the probability of acquiring a startup to firm sales. We estimate

$$D_{fst}^{SU} = \alpha_t + \alpha_s + \beta \ln\left(\text{sale}_{fst}\right) + \varepsilon_{fst}, \tag{33}$$

where $D_{f,t}^{SU}$ is a dummy equal to 1 if firm f of industry s buys a startup in year t, sale_{fst} is equal to the sales of firm f in year t, and α_t and α_s are a set of year and industry (3-digit NAICS) fixed effects.

Table 2 shows the results for this regression. Clearly, larger firms are more likely to acquire startups: a doubling of firm size (an increase of 0.69 log points) is correlated with an increase in the probability of startup acquisition of about 0.2 percentage points. This compares to an average acquisition frequency (computed as the ratio of all startup acquisitions in SDC Platinum to the number of all firm-years in Compustat) of 3.26%. We will target these numbers for the calibration of our model.

Table 2: Startup acquisition by firm size

	(1)	(2)
ln(sale)	0.00292***	0.00294***
	(0.000115)	(0.000116)
Year FE	Yes	No
Industry FE	Yes	No
Industry-Year FE	No	Yes
R^2	0.02	0.03
Observations	104,263	104,178

Notes: This table shows the estimates for β in equation (33). The dependent variable is a dummy for whether or not an incumbent acquires a startup in a given year.

Having established these stylized facts, we now turn to the central part of our empirical analysis, assessing the effect of acquisitions on the implementation of startup ideas.

3.3 The effect of acquisitions on the implementation of ideas

As we have shown in Section 2, one 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) or decrease it (if

incumbents are engaging in killer acquisitions). In this section, we try to assess the relative strength of these forces in the data.

To do so, we need a proxy for the implementation of startup ideas. We propose to rely on the evolution of patent citations. Precisely, 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.[25]

Of course, just considering the change in patent citations after acquisition faces an endogeneity problem, as acquired patents are different from the average patent (see Table 1). Therefore, we use a nearest neighbor matching algorithm to link each treated patent (belonging to an acquired startup) to up to ten control patents. We match on several patent and assignee characteristics, including patent application year, technological subsector, the target firm's founding year, or the number of inventors that develop each patent. 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.1 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:

$$\text{NumCites}_{ijt} = \beta_1 \text{Acquired}_{ij} + \beta_2 \text{Post}_{jt} + \beta_3 \Big(\text{Acquired}_{ij} \times \text{Post}_{jt} \Big) + \alpha_j + \alpha_t + u_{ijt}, \quad (34)$$

where NumCites_{ijt} is the number of citations received by patent i belonging to patent pair j at year t,[26] Acquired_{ij} is a dummy variable taking value 1 if patent i is acquired, and Post_{jt} is a dummy variable taking value 1 if year t is later than the acquisition of the acquired patent in pair j. In our baseline analysis, we consider a 10-year window around the acquisition (i.e., 5 years before and 5 years after this event). Finally, α_j are matched patent pair fixed effects and α_t are year fixed effects.

In equation (34), 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

[25] In line with this interpretation, Argente *et al.* (2020) show that in the consumer goods sector, more highly cited patents lead to a higher likelihood of introducing new products.

[26] A 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. The dependent variable is winsorized at the 95% level.

Poisson estimator without fixed effects. The interaction term is positive, but not statistically significantly different from zero. When adding matched patent pair fixed effects (column 2), year fixed effects (column 3), or both sets of fixed effects (column 4), the point estimate becomes even closer to zero. Thus, our results appear to indicate that on average, acquisitions have no effect on startup patent citations. 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. Appendix B.4 discusses several additional robustness checks, including changes in the number of control patents and different event study windows.

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

	(1)	(2)	(3)	(4)
Post	1.347***	1.395***	1.286***	0.721***
	(0.129)	(0.113)	(0.151)	(0.139)
Acquired	-0.316	-0.242	-0.298	-0.209
	(0.484)	(0.451)	(0.459)	(0.437)
Acquired x Post	0.328	0.297	0.318	0.262
	(0.557)	(0.540)	(0.544)	(0.525)
Matched Patents FE	No	Yes	No	Yes
Year FE	No	No	Yes	Yes
Observations	22,373	22,119	22,373	22,119

Notes: We use a Poisson estimator. The dependent variable is the number of citations received at the patent-year level. *Acquired* is a dummy equal to 1 for acquired patents, and *Post* is a dummy equal to 1 for post-acquisition years. Standard errors are clustered at the target firm level. * significant at 10%; ** significant at 5%; *** significant at 1%.

Heterogeneous effects Our model also predicts that the effect of an acquisition on a startup idea depends on the characteristics of the acquirer. Most importantly, acquisitions of a related startup should have a more negative effect on implementation than acquisitions of an unrelated startup.

To test this prediction, we define related acquisitions as acquisitions in which the target and the acquirer belong to the same NAICS 3-digit industry, and estimate a triple-difference model:

$$\begin{aligned}\text{NumCites}_{ijt} =& \beta_1 \text{Acquired}_{ij} + \beta_2 \text{Post}_{jt} + \beta_3 \text{Related}_{ij} + \beta_4 \text{Acquired}_{ij} \times \text{Post}_{jt} \\ &+ \beta_5 \text{Acquired}_{ij} \times \text{Related}_{ij} + \beta_6 \text{Post}_{jt} \times \text{Related}_{ij}\end{aligned}$$

$$+ \beta_7 \text{Acquired}_{ij} \times \text{Post}_{jt} \times \text{Related}_{ij} + \alpha_j + \alpha_t + u_{ijt}. \quad (35)$$

Table 4 summarizes our results, with different specifications again corresponding to different combinations of fixed effects. Throughout all specifications, we find that for related acquisitions, citations to the startup patents fall more after an acquisition than for unrelated acquisitions. This difference is statistically significant. Thus, the zero effect in the overall sample is the average of a negative effect of acquisitions on citations for related acquisitions, and a positive effect for unrelated acquisitions.

Furthermore, Table B.3 in Appendix B also shows that acquisitions in the pharmaceutical industry (for which Cunningham *et al.* (2021) found their "killer acquisitions" result) have on average a negative effect on citations, even though this effect is not statistically significant.

Table 4: Heterogeneous effects: related and unrelated acquisitions

	(1)	(2)	(3)	(4)
Post	1.329***	1.336***	1.267***	0.705***
	(0.170)	(0.134)	(0.198)	(0.155)
Acquired	-0.965**	-0.917**	-0.927*	-0.860*
	(0.414)	(0.459)	(0.478)	(0.474)
Acquired x Post	1.133**	1.099**	1.107*	1.037*
	(0.562)	(0.548)	(0.574)	(0.564)
Acquired x Related	1.127	1.175*	1.077	1.115*
	(0.699)	(0.642)	(0.699)	(0.621)
Post x Related	0.0383	0.139	0.0308	0.0388
	(0.211)	(0.171)	(0.218)	(0.174)
Acquired x Post x Related	-1.494*	-1.493*	-1.456*	-1.429*
	(0.804)	(0.782)	(0.805)	(0.757)
Matched Patents FE	No	Yes	No	Yes
Year FE	No	No	Yes	Yes
Observations	22,373	22,119	22,373	22,119

Notes: We use a Poisson estimator. The dependent variable is the number of citations received at the patent-year level. *Acquired* is a dummy equal to 1 for acquired patents, *Post* is a dummy equal to 1 for post-acquisition years, and *Related* is a dummy equal to 1 if the acquisition target and the acquirer belong to the same NAICS 3-digit industry. Standard errors are clustered at the target firm level. * significant at 10%; ** significant at 5%; *** significant at 1%.

Summing up, our results in this section suggest that a ban on startup acquisitions would not have a significant effect on the implementation of startup ideas (although a

ban limited to related acquisitions might have). However, this preliminary conclusion might well be misleading. First, our difference-in-differences regressions miss general equilibrium effects that affect all firms equally. Second, as our discussion in Section 2 shows, the implementation channel is not the only link between acquisitions and innovation: acquisitions also affect incumbent innovation and startup creation. In the next section, we return to a calibrated version of our model to jointly evaluate these forces. In doing so, our results from this section will be important calibration targets.

4 Calibration

4.1 Externally calibrated parameters

We set a number of parameters externally. We assume that in the baseline BGP, there are no acquisition taxes ($T_R = T_U = 0$), and that a period of unit length in the model corresponds to one year in the data. We set the discount rate to $\rho = 0.03$, which together with a 2% target for the annual growth rate implies a real interest rate of 5%. We assume that there are two levels of product spending shares, $\Omega = \{\omega_L, \omega_H\}$ with $\omega_L < \omega_H$. 10% of products belong to the H class, so $\tau(\omega_H) = 0.1$, and we set $\omega_H/\omega_L = 16.125$. These choices follow Hottman *et al.* (2016), who show that in a large sample of consumer goods the 90-th percentile of the product sales distribution is 16.125 times larger than the median. Finally, we set the cost elasticity of innovation and implementation to $\psi = \theta = 2$, following empirical evidence summarized in Akcigit and Kerr (2018).

4.2 Internally calibrated parameters

These choices leave ten parameters to be identified: the productivity step size, λ; the scale and curvature of the search technology of incumbents, χ and φ; the share of startup ideas matched to unrelated incumbents, γ; the incumbent's bargaining weight, α; the incumbent's internal and external innovation cost shifters, ξ_I^{int} and ξ_I^{ext}; the startup creation cost, ξ_S; and the implementation cost shifters for incumbents and startups, κ_I and κ_S.

We calibrate these parameters internally, using an indirect inference approach. That is, we choose values for the ten parameters that minimize the distance between ten model-generated moments and their empirical counterparts.[27] The success of this calibration strategy relies on choosing moments that are relevant for the economic mechanisms we

[27] Appendix C.2 describes how we compute all model moments. We then choose the vector of parameters $\boldsymbol{\theta} = (\lambda, \chi, \varphi, \gamma, \alpha, \xi_I^{\text{int}}, \xi_I^{\text{ext}}, \xi_S, \kappa_I, \kappa_S)$ to minimize the distance function $\sum_{m=1}^{10} \frac{|\text{Moment}_m(\text{Model},\boldsymbol{\theta}) - \text{Moment}_m(\text{Data})|}{0.5|\text{Moment}_m(\text{Model},\boldsymbol{\theta})| + 0.5|\text{Moment}_m(\text{Data})|}$.

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. Our discussion in the main text stresses the economic intuitions underlying identification. Appendix C.4 shows the results of a global identification test, indicating that all parameters are well identified by the chosen moments.

Startup acquisition moments We target several moments from our empirical analysis in Section 3. First, we aim to reproduce a share of unrelated acquisitions of 55%, corresponding to the share of startup acquisitions in which target and acquirer belong to different NAICS 3-digit industries (see Table 1). This moment pins down γ, the probability that a startup idea is matched to an unrelated incumbent.

Second, we target a startup acquisition frequency of 3.26% (the ratio of startup acquisitions to publicly listed firms, as reported in Section 3.2). Thus, we take publicly listed firms as the data equivalent of incumbents in our model. This is an approximation, and so this target is observed with some measurement error. However, Appendix C.5 shows that our results are very similar for other values of the acquisition frequency, ranging between 2 and 6%. This moment identifies χ, which shifts search costs and therefore governs the probability with which incumbents notice and acquire startups.[28]

Third, our regressions in Section 3.3 showed that the average acquisition has no effects on the citations of a startup patent. We replicate this regression in the model (using the implementation probability of a startup idea as the outcome variable), targeting a coefficient of zero.[29] This moment identifies the relative implementation cost of incumbents, κ_I/κ_S. Indeed, differences in implementation costs shift the relative implementation probability of incumbents with respect to startups. As shown in Table 5, our calibration implies that incumbents have about 32% lower implementation costs than startups. This compensates for related incumbents' lower implementation incentives, leading to the overall zero effect.

Fourth, we target an average acquisition premium of 49.3% (see Table 1). In the model, the acquisition premium is the percentage difference between the acquisition price and the startup's outside option. This moment pins down α, the incumbent's bargaining weight. We find that $\alpha = 0.75$, implying that incumbents obtain the largest part of the surplus.

Fifth, we target the fact that large firms are more likely to acquire startups, matching

[28] When solving our model, we only needed to characterize product-level (as opposed to firm-level) outcomes. However, to compute this moment, we need to know the mass of incumbent firms. To do so, we simulate a large cohort of firms, allowing us to compute the average number of products per firm (and the joint distribution of spending shares and technology gaps across firms). Appendix C.3 contains further details.

[29] As we argued earlier, there is a positive link between citations and implementation. However, as the targeted regression coefficient is zero, we do not need to assume an elasticity of implementation to citations.

the coefficient from our regression of an acquisition dummy on firm sales (see Table 2). We compute the model equivalent of this moment by running this regression in the model. The moment identifies the curvature φ in the search cost function, which governs how steeply search costs increase for firms that search harder.

Finally, we assume that startups which are not acquired on average implement their idea (and hence enter) with a probability of 10%. This target pins down the level of implementation costs κ_S. While this statistic is hard to measure, the database of Guzman and Stern (2020) indicates that 6.6% of non-acquired patenting startups either achieve an IPO or grow to more than 100 employees. Given the construction of their data (which links incorporation records to other datasets and might miss some matches), this number is arguably a lower bound. Thus, we choose a somewhat higher target of 10%. However, as we show in Appendix C.5, this target matters little for our quantitative results, which would be virtually unchanged with a target of 5% or 15%.[30]

Growth, growth contributions and exit rates We target a growth rate of 2%, the long-run average growth rate of GDP per person in the United States (Jones, 2016). This identifies the innovation step λ which shifts growth up or down (see equation (30)).

Finally, we discipline the drivers of growth by following the influential work of Garcia-Macia *et al.* (2019), who use micro-level data from the Census Bureau's Longitudinal Business Database to structurally estimate the growth contributions of different firm and innovation types. They find that over the period 1993-2013, 21.1% of growth is due to entrants and 79.9% to incumbents. We target the same shares in the model. This identifies the level of incumbent innovation costs, $\tilde{\zeta}_I^{\text{int}}$ and $\tilde{\zeta}_I^{\text{ext}}$, which determines incumbents' growth contribution in our model. Garcia-Macia *et al.* (2019) also find that 12.6% of growth from incumbent innovation stems from creative destruction (the remainder being due to incumbents' improvement of their existing products). This moment identifies the relative cost of external innovations, $\tilde{\zeta}_I^{\text{ext}}/\tilde{\zeta}_I^{\text{int}}$. Finally, computing exit rates over five-year intervals, Garcia-Macia *et al.* (2019) find an exit rate of 6% for large firms (i.e., firms with above-average employment) between 1993 and 2013. We target this number as the exit rate for incumbents in our model.[31] This identifies the startup creation cost $\tilde{\zeta}_S$, which through the

[30] Roughly speaking, growth depends on the own innovation rate of startups, which is pinned down by our target for the exit rate. 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.

[31] This is an approximation, as large firms in Garcia-Macia *et al.* (2019) do not correspond to publicly listed firms in the data. However, using other studies would imply a similar target. For instance, structurally estimating an endogenous growth model, Akcigit and Kerr (2018) find an entry rate for innovative firms of 5.8%. This is in line with our target, as entry and exit rates coincide in our model.

free entry condition shifts the startup rate and hence the exit rate of incumbents.

Table 5: Parameter values and model fit

Parameter	Value	Interpretation	Source / Moment	Data	Model
ρ	0.03	Discount rate	Standard value		
τ_H	0.1	Fraction of H products	Normalization		
ω_H/ω_L	16.125	Rel. spending share, H products	Hottman *et al.* (2016)		
ψ	2	Incumbent innovation elasticity	Akcigit and Kerr (2018)		
θ	2	Implementation elasticity	Akcigit and Kerr (2018)		
λ	1.041	Innovation step size	Growth rate	2%	2%
χ	0.279	Search cost scale	Frequency of startup acq.	0.0326	0.0332
κ_I/κ_S	0.681	Inc. relative impl. cost	Reg., implem. on acq.	0	-0.004
α	0.745	Inc. bargaining weight	Acquisition premium	49.3%	49.4%
φ	2.615	Search cost curvature	Reg., acq. on firm sales	0.0029	0.0029
γ	0.639	Share of unrelated startups	Share of unrelated acq.	54.9%	55%
ξ_I^{int}	0.918	Inc. int. innovation cost	Entry contrib. to growth	21.1%	20.7%
$\xi_I^{\text{ext}}/\xi_I^{\text{int}}$	5.715	Rel. ext. innovation cost	Incumbent CD share	12.6%	13%
ξ_S	0.061	Startup cost	Exit rate	6%	6%
κ_S	2.317	Startup implementation cost	Average impl probability	10%	10%

Notes: This table lists the parameter values for our baseline calibration. For each internally calibrated parameter, the table lists the moment identifying the parameter, as well as the data and model value of that moment.

Model fit and untargeted moments Table 5 lists the calibrated parameter values and shows that our model fits the ten targeted moments closely. This is due to the tight link between targeted moments and parameters, shown more formally in Appendix C.4.

Our model is also in line with a number of untargeted moments. Figure 5 compares the sales distribution of incumbent firms in the model to the sales distribution of publicly listed firms in Compustat, by plotting the Lorenz curve of firm sales for both model and data.[32] The firm sales distribution in our model is strongly skewed, even though it is not quite as skewed as the one in the data.

Our calibrated model is also qualitatively in line with our regression results shown in Table 4, indicating that startup patent citations fall after related acquisitions and increase after unrelated acquisitions.[33]

[32] In the data, we first compute the Lorenz curve for each 3-digit NAICS industry in the year 2010. That is, we compute the share of total industry sales for each decile of the industry sales distribution. We then take a weighted average of these values, where the weights are total industry sales. In the model, we compute the sales distribution of firms using the same model-generated data that we use to compute our targeted moments. See Appendix C.3 for details.

[33] Quantitatively, our model predicts a difference in implementation probabilities of 0.87 log points. In the data, we find a difference of 1.06 log points for patent citations. These numbers are not directly comparable: to confront them, we need to know the elasticity of the implementation probability of ideas (the outcome in our model) to patent citations (the outcome in our empirical analysis). While this elasticity is hard to measure in the data, Kogan *et al.* (2017) estimate that the elasticity of a patent's market value to its number of forward citations is 0.17. In our model, the expected value of an idea is equal to the product of the implementation

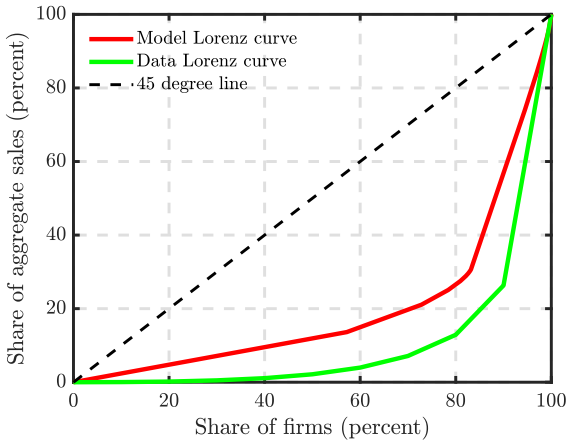

Figure 5: The Lorenz curve of firm sales in the model and in the data

5 The aggregate effects of startup acquisition bans

We are now ready to tackle the central question of our paper: how do limits to startup acquisitions affect innovation, growth and welfare?

5.1 Startup acquisition bans: partial and general equilibrium effects

In this section, we focus on the simplest possible policy, a ban on all startup acquisitions. Throughout, we rely on our insights from Section 2.4 to decompose the growth effect of this policy into three margins: changes in the startup rate, changes in incumbent innovation, and changes in the implementation rate of startup ideas.

Partial equilibrium effects The policy discussion around killer acquisitions has focused on the third margin in our decomposition: the implementation rate of startup ideas. Intuitively, if acquisitions lower the implementation probability of startup ideas, banning them should increase growth. On the other hand, if acquisitions increase the implementation probability of startup ideas, banning them should lower growth.

Columns (B) and (C) in Table 6 assess this effect in isolation. That is, they compute growth assuming that no startups are acquired, but all other decisions (i.e., the startup rate, incumbent innovation policies, and startup implementation rates) are unchanged.[34] This is

probability and the value of the implemented idea, and the latter is ex ante identical for all startup ideas. Using this elasticity, the data estimate for citations translates to a difference in implementation probabilities of 0.26, somewhat lower than what our model predicts.

[34]Formally, we compute growth by setting the noticing probabilities s_U and s_R in equation (30) to zero, leaving the policy functions i_S, x_S, x and z, as well as the distribution m unchanged.

obviously not an equilibrium outcome of our model, but it is a useful partial equilibrium exercise to zoom in on the role of the implementation rate. As shown in Column (C), this partial equilibrium assessment of the acquisition ban suggests an increase in startup idea implementation, growth and welfare. The growth rate increases by 0.9% (from 2.00% to 2.02%), and welfare increases by 1.0%.

Table 6: The effects of a startup acquisition ban.

Outcome	(A) Baseline	(B) PE Acq Ban	(C) Change	(D) GE Acq Ban	(E) Change
Growth rate	2%	2.02%	0.9%	1.97%	-1.6%
Acquisition frequency	0.033	0	-100%	0	-100%
Incumbent own innov.	0.381	0.381	0%	0.382	0.4%
Startup rate	0.403	0.403	0%	0.335	-16.8%
Startup idea implem.	0.294	0.305	3.8%	0.326	11%
Welfare			1%		-0.3%

Notes: This table summarizes the effect of a ban on all startup acquisitions. Column (A) shows some summary statistics for our baseline calibration. Column (B) shows the partial equilibrium implications of an acquisition ban, re-computing key outcomes by setting the noticing probabilities s_U and s_R to 0, but leaving all other endogenous variables unchanged. Column (D) shows the general equilibrium implications of an acquisition ban, by computing the BGP equilibrium in the case in which acquisition taxes T_R and T_U are high enough to eliminate all acquisitions. Column (C) shows the percentage change between Column (A) and (B), and Column (E) the percentage change between Column (A) and Column (D).

This result is surprising: our empirical results showed that the average acquisition had no effect on implementation, and our model matches this result. Given this, how can an acquisition ban increase implementation? The answer to this question lies in the heterogeneity of acquisitions. Unrelated acquisitions (which increase implementation) are equally likely for low and high spending share products. However, related acquisitions (which lower implementation) are more likely for high spending share products, as shown in Figure 2. As innovation on high spending share products contributes more to growth, the positive effect of banning related acquisitions dominates. Finally, welfare increases due to higher growth, and as resources used for startup search can now be used for consumption.[35]

Summing up, with unchanged startup rates and incumbent innovation, a startup acquisition ban would increase growth. However, this assessment is obviously incomplete: in equilibrium, startup rates and incumbent innovation adjust to an acquisition ban. Quantifying these responses is the main objective of our paper, and we turn to this next.

General equilibrium feedback Columns (D) and (E) of Table 6 summarize the general equilibrium effects of an acquisition ban. To do so, we compare our baseline BGP to

[35]The representative consumer's consumption-equivalent welfare change from going from a BGP A to another BGP B is given by $\frac{C_0^B}{C_0^A}\exp\left(\frac{g_B-g_A}{\rho}\right)$.

an alternative BGP in which the government eliminates startup acquisitions by imposing arbitrarily high acquisition taxes T_R and T_U.[36] We find that general equilibrium forces reverse the positive partial equilibrium effects of the acquisition ban: growth decreases by 1.6% (from 2.00% to 1.97%) and welfare falls slightly, by 0.3%.

The small change in the growth rate is driven by large changes in the underlying margins. As a result of the ban, the startup rate falls by 16.8%.[37] This decrease is only partially offset by a 0.4% increase in the own innovation of incumbents and an 11% increase in the implementation rate of startup ideas.

To understand these effects, consider first the startup rate. Startups agree to acquisitions if the acquisition price exceeds their outside option. Thus, everything else equal, acquisitions create surplus value for them. Banning acquisitions eliminates this surplus value, and triggers a fall in the equilibrium startup rate. This lower startup rate has a direct negative effect on growth, but it also has wide-ranging spillover effects for incumbents.

The effect of the ban on incumbents is ambiguous. On the one hand, there is a direct negative effect: the ability to acquire related and unrelated startups also created a surplus for incumbents. A ban eliminates this surplus, and hence all else equal lowers incumbent value and innovation incentives. On the other hand, the lower startup rate has a positive spillover effect: it implies that incumbents are less likely to be displaced or threatened by startups, thus increasing their value and innovation incentives. Overall, this positive effect dominates, and average incumbent value increases. However, effects are heterogeneous across the technology gap distribution. The positive effect of startup bans is concentrated among products with low technology gaps (which are most vulnerable to displacement), while the negative effect is concentrated among products with high technology gaps (which have a higher surplus from acquisitions). Thus, as shown in the left panel of Figure 6, the acquisition ban flattens the value function. As a result, the right panel of Figure 6 shows that incumbent own innovation slightly falls for most products, while incumbent external innovation (which depends on average value) increases. Finally, the underlying technology gap distribution also changes as a result of the ban, shifting to the left (i.e., to lower technology gaps). All else equal, this also increases incumbent innovation. The net effect of all these forces is a small increase in incumbent innovation.

The increase in average incumbent value also implies an increase in the implementation incentives of startups (which invest in implementation in order to become incumbents).

[36] In practice, we set taxes to a level ensuring that the surplus from startup acquisitions is always zero. Thus, no incumbent pays the tax in equilibrium, and there is no tax revenue for the government.

[37] This response is within the range of estimates computed by Eisfeld (2023) for the enterprise software market. Eisfeld estimates that if all startup acquisitions were blocked, startup creation would fall by 8-20%.

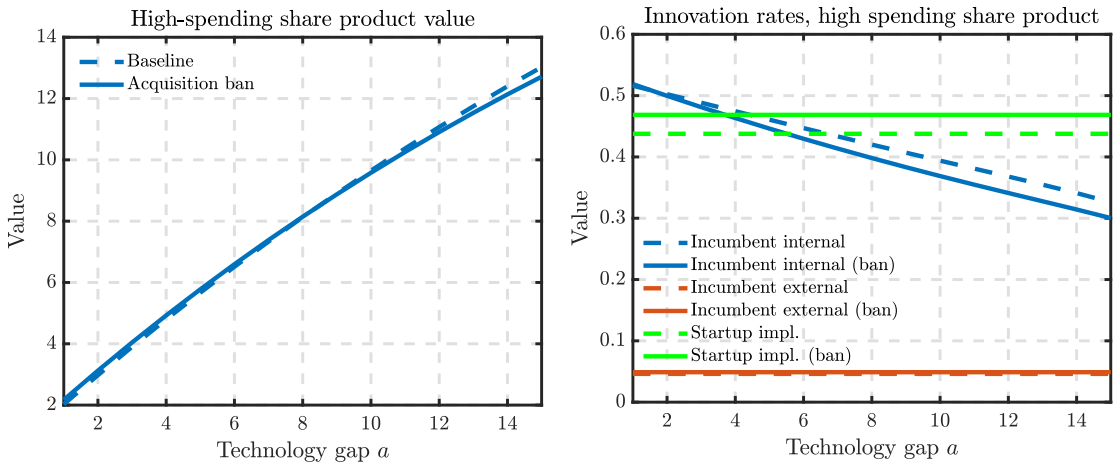

Figure 6: General equilibrium effects of an acquisition ban

Therefore, the general equilibrium response of the implementation rate of startup ideas is substantially larger than the partial equilibrium one, as shown in Table 6.

Finally, the welfare effects of the startup acquisition ban are less negative than the growth effects, with welfare decreasing by just 0.3%. This is due to the fact that the policy reduces both search and innovation spending (especially for startup creation), hence increasing the resources available for consumption.

This concludes our discussion of overall bans. In the next section, we turn to more sophisticated policies, which selectively ban different types of acquisitions.

5.2 Banning certain types of acquisitions

Table 7 summarizes the results of policies that ban only related acquisitions, or only unrelated acquisitions. As before, we start by considering only the partial equilibrium effects on the implementation rate of startup ideas. Then, the effects of partial bans are straightforward. Related acquisitions are on average killer acquisitions. Hence, banning them increases the implementation rate of startup ideas, growth and welfare, as shown in Column (A).[38] Unrelated acquisitions, in turn, increase the implementation rate of startup ideas, and banning them lowers growth, as shown in Column (C).

However, taking into account general equilibrium forces reverses these effects. As Column (B) shows, a ban on related acquisitions strongly lowers the startup rate, leading to lower growth and welfare. The mechanisms (and indeed their quantitative magnitude) are similar to the effects of an overall ban. Indeed, all of the displacement risk and most of

[38]This policy is roughly in line with the Platform Competition and Opportunity Act debated by Congress, which would prohibit all acquisitions of direct, nascent or potential competitors (for Technology Platforms).

Table 7: Banning certain types of acquisitions

Outcome	(A) PE Rel. Ban	(B) GE Rel. Ban	(C) PE Unrel. Ban	(D) GE Unrel. Ban
Growth rate	1.2%	-1.6%	-0.3%	0.1%
Acquisition frequency	-45%	-53.2%	-55%	-54.7%
Incumbent own innov.	0%	0.3%	0%	0.1%
Startup rate	0%	-17.5%	0%	1.3%
Startup idea implem.	5.2%	12.2%	-1.4%	-1.4%
Welfare	1.2%	-0.3%	0.1%	0%

Notes: This table summarizes the effect of bans on related and unrelated startup acquisitions. Columns (A) - (D) show changes in key outcomes with respect to the baseline BGP. Column (A) shows the partial equilibrium implications of a ban on related acquisitions, re-computing key outcomes by setting the noticing probabilities s_R to 0, but leaving all other endogenous variables unchanged. Column (B) shows the general equilibrium implications of a ban on related acquisitions, by computing the BGP equilibrium in the case in which the government imposes a tax T_R that is high enough to eliminate all related acquisitions. Column (C) and Column (D) are the equivalents of Columns (A) and (B) for unrelated acquisitions.

the surplus from acquisitions is due to related acquisitions. Thus, related acquisitions were already driving the bulk of the general equilibrium effects of the overall ban.

On the other hand, unrelated acquisition bans generate only a small general equilibrium feedback, owing to their small surplus. Banning these acquisitions raises the startup rate. This is due to the fact that unrelated acquisitions increase creative destruction: some startup ideas are taken over by unrelated incumbents that are better suited to implement them. When this is shut down, incumbents are less likely to be displaced, which increases their value and hence the startup rate and incumbent own innovation. These positive effects almost exactly cancel out the negative partial equilibrium effects from lower implementation, so that the net effect of the ban is essentially zero.

Summing up, our results highlight the importance of general equilibrium effects for assessing limits to startup acquisitions. In the next section, we discuss how the magnitude of these general equilibrium effects depends on our calibration choices.

5.3 Do startup acquisition bans always lower growth?

The sign and magnitude of our quantitative results are not hard-wired features of our model. Instead, they depend on our calibration targets. In this section, we systematically explore the role of these targets, and also consider a more general formulation of the free entry assumption. This gives a sense of the robustness of our results, and also indicates circumstances in which acquisition bans could be more or less desirable.

Alternative calibration targets Figure 7 illustrates the role of what turns out to be our most important calibration target, the acquisition premium. The baseline value for this

target is 49.3%, indicated by the vertical lines. Each point on the x-axis shows the effects of an acquisition ban for a re-calibrated model with a different value for the acquisition premium, leaving all other targets and all external parameters unchanged. The left panel shows changes in growth and welfare, while the right panel decomposes the growth effect into the familiar three margins.

As the figure shows, the acquisition premium target is crucial for our results. Higher acquisition premia make bans more desirable: if the premium exceeds 60%, bans increase welfare, and if it exceeds 90%, bans also increase growth. These results are due to the fact that the acquisition premium disciplines the costs and benefits of acquisitions. High premia imply that acquisitions are lucrative for startups, but costly for incumbents. Thus, with high premia, banning acquisitions has a more negative effect on the startup rate (a highly lucrative option disappears), but a more positive effect on incumbent innovation (incumbents lose little by renouncing expensive acquisitions, and gain from a large decrease of creative destruction by entrants). As incumbents contribute more to aggregate growth than startups, the latter effect dominates.

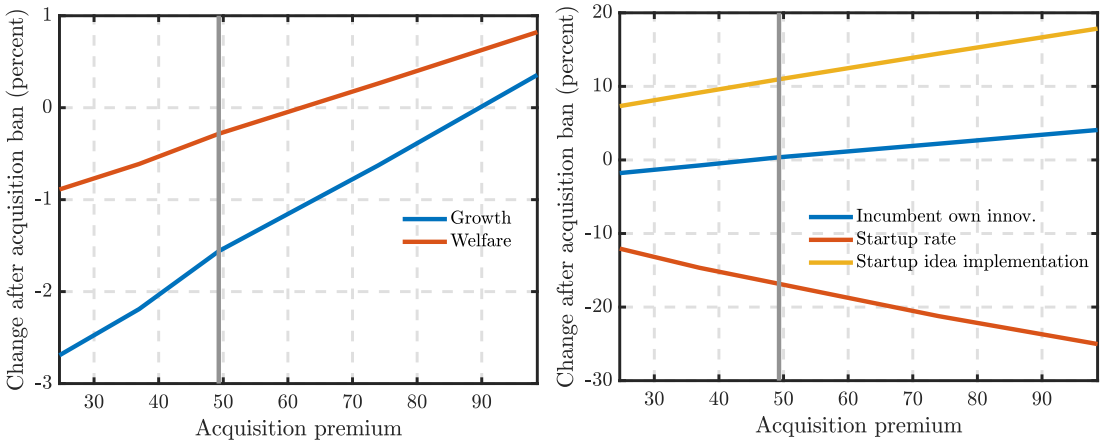

Figure 7: Robustness: Different targets for the acquisition premium. *Notes:* Each point on the x-axis corresponds to a re-calibrated model, leaving all other calibration targets and external parameters unchanged.

These results imply that startup acquisition bans might be more effective in environments with high acquisition premia. For instance, our data indicates that the average acquisition premium for a target in the pharmaceutical industry (NAICS Code 3254) is 95%. Although this number is based on a very small sample and therefore needs to be interpreted with caution,[39] this suggests that startup acquisitions bans might be more desirable in the

[39]There are 13 targets from the pharmaceutical industry that have information on the acquisition premium. Conversely, we have also 13 observations for acquirers from the pharmaceutical industry. The average acquisition premium in this group is 97%.

pharmaceutical industry than elsewhere.

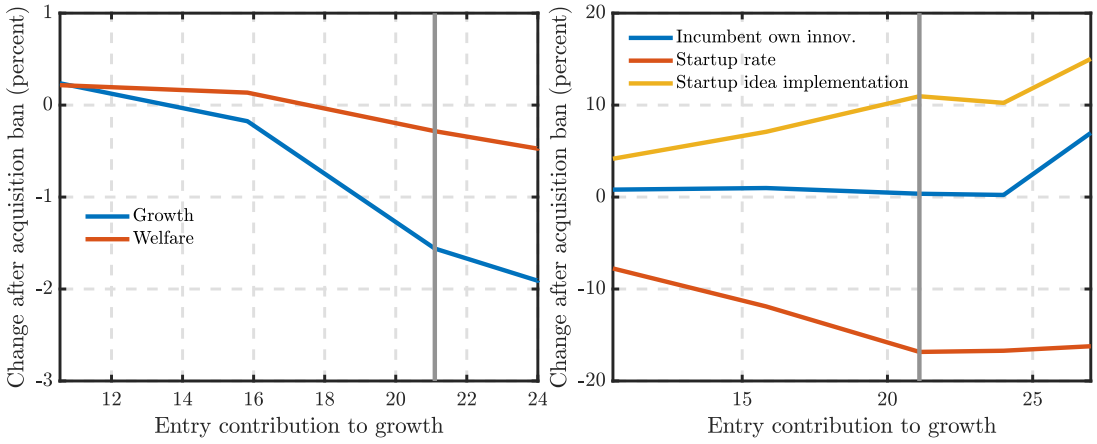

Figure 8: Robustness: different targets for the growth contribution of entrants. *Notes:* See Figure 7.

Figure 8 conducts the same exercise by considering different targets for the contribution of entrants to growth. As the left panel shows, a lower target dampens the negative growth effect of an acquisition ban. Indeed, when entrants contribute less to growth, the fall in the startup rate triggered by the ban matters less for aggregate growth, and the increase in incumbent innovation matters more.

Finally, Figure 9 examines the role of our empirical estimate for the effect of acquisitions on the implementation probability of startup ideas. The baseline target for this number was zero. Unsurprisingly, targeting instead a positive effect means that an acquisition ban lowers growth even more: for instance, targeting an increase of 0.1 log points would imply that a ban lowers growth by 2.1% rather than by 1.6%. Indeed, as the right panel shows, changes in this moment directly move the implementation rate of startup ideas (while the other margins remain roughly unchanged). However, the figure also shows that acquisitions would need to have a very large effect on implementation to move the needle. With an elasticity of implementation to citations of 0.17 (as discussed in footnote 33), an implementation increase of 0.1 log points would correspond to a 0.6 log points increase in citations, which is substantially larger than the results found in our regressions.

Alternative free entry assumptions Finally, given the important role of the startup rate for our results, it is useful to discuss our free entry assumption. Our model has a linear entry cost. While this is standard in the literature (Klette and Kortum, 2004; Peters, 2020), it obviously influences the strength of the response of the startup rate to an acquisition ban.

To assess the robustness of our results, we now relax this assumption. Precisely, we assume that startup creation costs are given by $\xi_S (x_S)^\nu Y_t$. In our baseline model, $\nu = 0$.

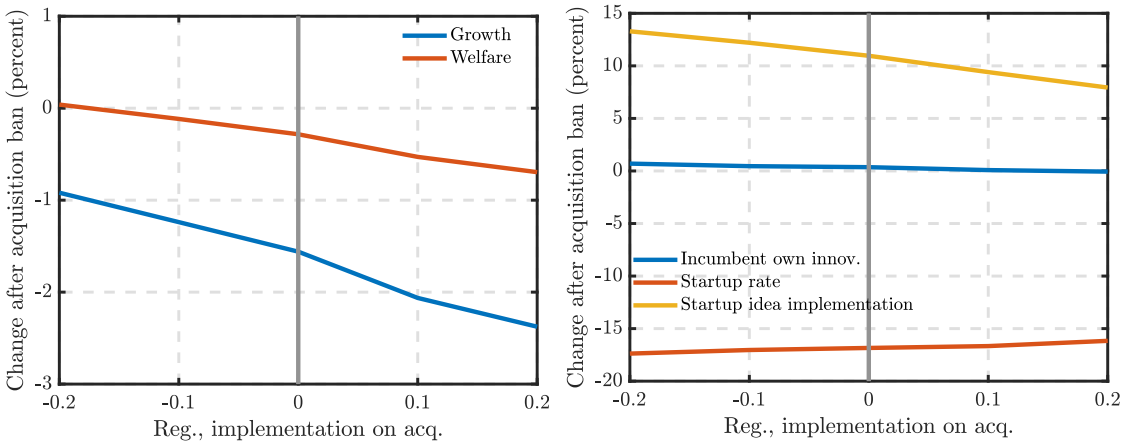

Figure 9: Robustness: different targets for the causal effect of acquisitions on the number of citations of startup patents. *Notes:* See Figure 7.

Instead, allowing ν to be a positive parameter generates congestion effects in startup creation, with the cost of startup creation increasing in the mass of startups. We consider values of $\nu = 0.14$ (in line with the estimation results of Klenow and Li, 2024) and $\nu = 1$. We then re-calibrate the model to match the baseline targets, leaving all other external parameters unchanged.

Table 8 shows how congestion effects affect our conclusions. As expected, a higher congestion externality limits the fall in the startup rate after an acquisition ban. However, the overall decrease in the growth rate after a ban becomes even more negative. Indeed, as the startup rate decreases less, incumbents suffer the cost of disappearing acquisition surpluses without getting as much of the benefits of lower creative destruction by entrants. Thus, incumbent own innovation now decreases after a ban. As in our baseline analysis, the effects of overall bans and selective bans on related acquisitions are very similar.

Table 8: Robustness checks: external parameters

Parameter values	Growth	Inc. own inn.	Startup rate	Impl. rate su ideas
$\nu = 0.14$, Overall ban	-1.7%	-0.2%	-15%	9.6%
$\nu = 0.14$, Rel. ban	-1.7%	-0.3%	-15.6%	10.9%
$\nu = 1$, Overall ban	-2%	-1.6%	-8.4%	5.7%
$\nu = 1$, Rel. ban	-2%	-1.8%	-8.8%	6.7%

Notes: This table summarizes the robustness of our results to different choices for the congestion parameter ν. Each row corresponds to a different calibration, and in each calibration, all targeted moments and other external parameters are the same as in the baseline. Each column gives the percentage change of a key variable after an acquisition ban.

Overall, this section suggests that our quantitative results are robust to reasonable

variations in our calibration targets and model assumptions. However, they also indicate that the effects of startup acquisitions depend on the characteristics of the economy (and could therefore vary across different countries or industries). In particular, acquisitions are more problematic when entrants account for a small share of growth, when acquisitions decisively lower implementation probabilities, and when acquisition premia are high. In Appendix C.5, we expand on this analysis by conducting further robustness checks with respect to all calibration targets.

6 Conclusion

Our paper assesses the effect of startup acquisitions on aggregate growth. We use a general equilibrium model that takes into account a large number of potential positive and negative effects of these operations, and discipline it by calibrating it to micro-level data. Our results indicate that a startup acquisition ban would lower the startup rate, by reducing incentives for startup creation. This is only partially compensated by an increase in incumbents' own innovation and in the implementation rate of startup ideas. Accordingly, the ban lowers the growth rate by around 0.03 percentage points per year, and consumption-equivalent welfare by 0.3%.

As we have shown in the last section of the paper, our results are driven by the data used to discipline the model. Therefore, the effects of acquisition bans could vary depending on the country and time period considered, and might also be heterogeneous across industries. For instance, startup acquisitions are more problematic when incumbents implement startup ideas at low rates, account for a large share of growth, and pay high acquisition premia. Further exploring this heterogeneity is a promising path for future research.

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.

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. (2022). M&A and Early Investment Decisions by Digital Platforms. *Economia e Politica Industriale: Journal of Industrial and Business Economics*, **49** (3), 509–543.

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. (2022). The Novelty of Innovation: Competition, Disruption, and Antitrust Policy. *Management Science*, **68** (1), 37–51.

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., PETERS, R. and WANG, S. (2023). Measuring intangible capital with market prices. *Management Science (Forthcoming)*.

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.

HOTTMAN, C., REDDING, S. J. and WEINSTEIN, D. E. (2016). Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics*, **131** (3), 1291 – 1364.

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. (2022). Startup Acquisitions, Relocation, and Employee Entrepreneurship. *Strategic Management Journal*, **43** (11), 2189–2216.

KLENOW, P. J. and LI, H. (2024). *Entry Costs Rise with Growth*. 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.

OLMSTEAD-RUMSEY, J., PUGLISI, F. and WU, L. (2024). Digital Platform Acquisitions and Growth. *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.

WASI, N. and FLAAEN, A. (2015). Record Linkage Using Stata: Preprocessing, Linking, and Reviewing Utilities. *The Stata Journal*, **15** (3), 672–697.

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 Effects of Startup 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 The product-level value function

Consider an incumbent firm that produces $n = 1, 2, 3, \ldots$ products, and denote by $\mathbf{n} \equiv \{(\omega_j, a_j)\}_{j=1}^{n}$ the multiset of spending shares and technology gaps of each of the products of this firm. With our assumptions, the value function of an incumbent firm solves the following HJB equation:[40]

$$
\begin{aligned}
r_t V_t(\mathbf{n}, n) = \max_{\substack{\{z_j, s_{R,j}\}_{j=1}^{n};\\ s_U;\ x}} \Bigg\{ & \sum_{j=1}^{n} \Bigg[\omega_j \left(1 - \frac{1}{\mu(a_j)}\right) Y_t - \zeta_I^{\text{int}} z_j^{\psi} Y_t - \chi\left(s_{R,j}^{\varphi} + s_U^{\varphi}\right) Y_t \\
& + z_j \Big(V_t\Big(\mathbf{n} \backslash \{(\omega_j, a_j)\} \cup \{(\omega_j, a_j + 1)\}, n\Big) - V_t(\mathbf{n}, n) \Big) \\
& + x_{S,t} \Big(\alpha(1-\gamma) s_{R,j} \max\Big(B_{R,jt}(\mathbf{n}, n), 0\Big) + (1 - d_{jt}) V_t\Big(\mathbf{n} \backslash \{(\omega_j, a_j)\}, n-1\Big) - V_t(\mathbf{n}, n) \Big) \\
& + \widetilde{x}_t \Big(V_t\Big(\mathbf{n} \backslash \{(\omega_j, a_j)\}, n-1\Big) - V_t(\mathbf{n}, n) \Big) \Bigg] + \gamma x_{S,t} n \alpha s_U \int_0^1 \max\Big(B_{U,j't}(\mathbf{n}, n), 0\Big) dj' \\
& - \zeta_I^{\text{ext}} n x^{\psi} Y_t + n x \left(\int_0^1 V_t\Big(\mathbf{n} \cup \{(\omega_{j'}, 1)\}, n+1\Big) \mathrm{d}j' - V_t(\mathbf{n}, n) \right) \Bigg\} + \dot{V}_t(\mathbf{n}, n). \qquad \text{(A.1)}
\end{aligned}
$$

In this equation, we have defined the following variables. First, $B_{R,jt}(\mathbf{n}, n)$ denotes the value of the acquisition surplus when the incumbent producer of product j buys a startup with an idea on that product. This is given by

$$
B_{R,jt}(\mathbf{n}, n) \equiv \max_{i_{R,j}} \left\{ i_{R,j} \left(V_t\Big(\mathbf{n} \backslash \{(\omega_j, a_j)\} \cup \{(\omega_j, a_j + 1)\}, n\Big) - V_t(\mathbf{n}, n) \right) - \kappa_I i_{R,j}^{\theta} Y_t \right\}
$$

[40] In this equation, $\cup$ and $\backslash$ are multiset union and difference operators, i.e. for any two distinct elements a and b, we have $\{a, b\} \cup \{b\} = \{a, b, b\}$ instead of $\{a, b\} \cup \{b\} = \{a, b\}$, and $\{a, b, b\} \backslash \{b\} = \{a, b\}$ instead of $\{a, b, b\} \backslash \{b\} = \{a\}$.

$$- T_R Y_t - \Big(- i_{S,jt} V_t(\mathbf{n}, n) \Big) - \max_{i_{S,j}} \left\{ i_{S,j} V_t\Big(\{(\omega_j, 1)\}, 1 \Big) - \kappa_S i_{S,j}^{\psi} Y_t \right\}. \quad \text{(A.2)}$$

Second, $B_{U,j't}(\mathbf{n}, n)$ denotes the value of the acquisition surplus for an unrelated startup idea, i.e. when the incumbent acquires a startup with an idea on a product j' that the incumbent currently does not produce. This is given by

$$\begin{aligned} B_{U,j't}(\mathbf{n}, n) \equiv \max_{i_{U,j'}} \left\{ i_{U,j'} \Big(V_t\Big(\mathbf{n} \cup \{(\omega_{j'}, 1)\}, n+1 \Big) - V_t(\mathbf{n}, n) \Big) - \kappa_I i_{U,j'}^{\theta} Y_t \right\} - T_U Y_t \\ - \max_{i_{S,j'}} \left\{ i_{S,j'} V_t\Big(\{(\omega_{j'}, 1)\}, 1 \Big) - \kappa_S i_{S,j'}^{\psi} Y_t \right\}. \end{aligned} \quad \text{(A.3)}$$

Third, d_{jt} denotes the probability that the incumbent is displaced by a startup, given by:

$$d_{jt} = \gamma s_{U,t} i_{U,jt} + (1 - \gamma s_{U,t}) i_{S,jt}, \quad \text{(A.4)}$$

where $s_{U,t}$, $i_{U,jt}$, and $i_{S,jt}$ are policy functions (the solutions to the respective optimization problems written above). Finally, $x_{S,t}$ denotes the aggregate startup rate, and $\widetilde{x}_t$ is the aggregate creative destruction rate by (other) incumbents. The incumbent takes both of these rates as given.

Because net gains from external innovation and from the acquisition of unrelated startup ideas scale linearly in n, and the remaining terms are additively separable in j, it is easy to see that there exists a product-specific value function $\widehat{V}_t(\omega_j, a_j)$ for each $j \in \{1, \ldots, n\}$ so that $V_t(\mathbf{n}, n) = \sum_{j=1}^{n} \widehat{V}_t(\omega_j, a_j)$. The per-product value function solves the HJB equation:

$$\begin{aligned} r_t \widehat{V}_t(\omega_j, a_j) - \dot{\widehat{V}}_t(\omega_j, a_j) = \max_{z_j, s_{R,j}, s_U, x} \Bigg\{ & \omega_j \left(1 - \frac{1}{\mu(a_j)} \right) Y_t - \xi_I^{\text{int}} z_j^{\psi} Y_t - \chi\Big(s_{R,j}^{\varphi} + s_U^{\varphi} \Big) Y_t \\ & + z_j \Big(\widehat{V}_t(\omega_j, a_j + 1) - \widehat{V}_t(\omega_j, a_j) \Big) + x_{S,t} \Big(\alpha(1-\gamma) s_{R,j} \max\Big(\widehat{B}_{R,t}(\omega_j, a_j), 0 \Big) - d_{jt} \widehat{V}_t(\omega_j, a_j) \Big) \\ & - \widetilde{x}_t \widehat{V}_t(\omega_j, a_j) + \gamma x_{S,t} \alpha s_U \int_0^1 \max\Big(\widehat{B}_{U,t}(\omega_{j'}), 0 \Big) dj' - \xi_I^{\text{ext}} x^{\psi} Y_t + x \int_0^1 \widehat{V}_t(\omega_{j'}, 1) dj' \mathrm{d}j \Bigg\}, \end{aligned}$$

where $\widehat{B}_{U,t}$ and $\widehat{B}_{R,t}$ are now the per-product value of the acquisition surplus for unrelated and related startup ideas, respectively, equal to:

$$\widehat{B}_{R,t}(\omega_j, a_j) = \max_{i_{R,j}} \left\{ i_{R,j} \Big(\widehat{V}_t(\omega_j, a_j + 1) - \widehat{V}_t(\omega_j, a_j) \Big) - \kappa_I i_{R,j}^{\theta} Y_t \right\} - T_R Y_t$$

$$+ i_{S,jt}\widehat{V}_t(\omega_j,a_j) - \max_{i_{S,j}}\left\{i_{S,j}\widehat{V}_t(\omega_j,1) - \kappa_S i_{S,j}^{\psi}Y_t\right\}$$

$$\widehat{B}_{U,t}(\omega_{j'}) = \max_{i_{U,j'}}\left\{i_{U,j'}\widehat{V}_t(\omega_{j'},1) - \kappa_I i_{U,j'}^{\theta}Y_t\right\} - \max_{i_{S,j'}}\left\{i_{S,j'}\widehat{V}_t(\omega_{j'},1) - \kappa_S i_{S,j'}^{\psi}Y_t\right\} - T_U Y_t.$$

To solve the problem, we guess-and-verify that there exist functions $v^{\text{int}}:\Omega\times\mathbb{N}\to\mathbb{R}$, $v^{\text{ext}}:\Omega\to\mathbb{R}$, $b_R:\Omega\times\mathbb{N}\to\mathbb{R}$ and $b_U:\Omega\to\mathbb{R}$, such that:

$$\begin{aligned}
\widehat{V}_t(\omega_j,a_j) &= \overbrace{\left(v^{\text{int}}(\omega_j,a_j)+v^{\text{ext}}(\omega_j)\right)}^{\equiv\, v(\omega_j,a_j)}\, Y_t,\\
\widehat{B}_{R,t}(\omega_j,a_j) &= b_R(\omega_j,a_j)\, Y_t,\\
\widehat{B}_{U,t}(\omega) &= b_U(\omega_j)\, Y_t.
\end{aligned}$$

Our guess for the value of holding a product has two additive components. First, the firm receives value $v^{\text{int}}(\omega_j,a_j)$ from the discounted value of profits, internal innovation and related acquisition option values, which are a function of the product's spending share and its technology gap. Additionally, the firm perceives a franchise value $v^{\text{ext}}(\omega_j)$ coming from the option value of external innovations and unrelated startup acquisitions, both of which increase the value of the firm by expanding its number of products.

On the BGP, $\dot{\widehat{V}}_t(\omega_j,a_j) = v(\omega_j,a_j)gY_t$, where g is the aggregate rate of growth. Moreover, $\widetilde{x}_t=\widetilde{x}$, $x_{S,t}=x_S$, and our guesses imply that all products j with the same (ω_j,a_j) have the same value, which allows us to drop the j subscript everywhere. Plugging the guess and using the Euler equation $r=\rho+g$, we obtain:

$$\begin{aligned}
\rho\Big(v^{\text{int}}(\omega,a)+v^{\text{ext}}(\omega)\Big) = \max_{z,s_R,s_U,x}\Bigg\{ & \pi(\omega,a) - \zeta_I^{\text{int}} z^{\psi} - \chi\left(s_R^{\varphi}+s_U^{\varphi}\right)\\
& + z\Big(v^{\text{int}}(\omega,a+1)-v^{\text{int}}(\omega,a)\Big) - \widetilde{x}\Big(v^{\text{int}}(\omega,a)+v^{\text{ext}}(\omega)\Big)\\
& + x_S\Big(\alpha(1-\gamma)s_R\max\Big(b_R(\omega,a),0\Big) - d(\omega)\Big(v^{\text{int}}(\omega,a)+v^{\text{ext}}(\omega)\Big)\Big)\\
& - \zeta_I^{\text{ext}} x^{\psi} + x_S\gamma\alpha s_U\sum_{\omega'\in\Omega}\tau(\omega')\max\Big(b_U(\omega'),0\Big)\\
& + x\sum_{\omega'\in\Omega}\tau(\omega')\Big(v^{\text{int}}(\omega',1)+v^{\text{ext}}(\omega')\Big)\Bigg\},
\end{aligned} \tag{A.5}$$

where $\pi(\omega,a)=\omega\left(1-1/\mu(a)\right)$, $d(\omega)=\gamma s_U i_U(\omega)+(1-\gamma s_U)i_S(\omega)$. Taking first-order conditions yields equations (9), (10), (17) and (22) in the main text. We also recover the

expressions for $b_U(\omega, a)$ and $b_R(\omega)$ in equations (15) and (20).

Plugging back these results into equation (A.5), using the method of undetermined coefficients and the equilibrium condition $x = \widetilde{x}$, we find:

$$v^{\text{ext}}(\omega) = \frac{(\psi - 1)\,\xi_I^{\text{ext}} x^{\psi} + (\varphi - 1)\,\chi s_U^{\varphi}}{\rho + x + x_S d(\omega)}. \tag{A.6}$$

and

$$v^{\text{int}}(\omega, a) = \frac{\pi(\omega, a) + \xi_I^{\text{int}}\,(\psi - 1)\,z(\omega, a)^{\psi} + \chi\,(\varphi - 1)\,s_R(\omega, a)^{\varphi}}{\rho + x + x_S d(\omega)}. \tag{A.7}$$

This confirms our functional guesses for the value functions, with the value function in the main text holding $v(\omega, a) = v^{\text{int}}(\omega, a) + v^{\text{ext}}(\omega)$.

Using $(v^{\text{int}}, v^{\text{ext}})$, the ratio of implementation probabilities for a related startup idea (equation (24)) can be further decomposed as follows:

$$\left(\frac{i_R(\omega, a)}{i_S(\omega)}\right)^{\theta - 1} = \underbrace{\underbrace{\frac{v^{\text{int}}(\omega, a+1) - v^{\text{int}}(\omega, a)}{v^{\text{int}}(\omega, 1)}}_{\text{Displacement of product's profits}} \underbrace{\frac{v^{\text{int}}(\omega, 1)}{v^{\text{int}}(\omega, 1) + v^{\text{ext}}(\omega)}}_{\text{Displacement of firm's franchise value}}}_{\text{(a) Arrow replacement effect}} \underbrace{\frac{\kappa_S}{\kappa_I}}_{\text{(b) Cost differences}} \tag{A.8}$$

This equation decomposes the Arrow replacement effect into two components: one that operates through the profits of the preexisting product, and another one through the franchise value of the product. The startup has higher incentives to implement its idea both because it does not displace any profits from production and because it receives the franchise value, which the incumbent holds already.

A.2 The technology gap distribution

In this section, we characterize the technology gap distribution for each level of the spending share ω. Recall that spending shares are a fixed product characteristic (so that a product can never transition from one spending share to another). Therefore, the total mass of products with a given spending share, $\tau(\omega)$, is an exogenous parameter.

For any technology gap $a > 1$, we have

$$\begin{aligned}\Big(z(\omega, a-1) + x_S\,(1 - \gamma)\,s_R(\omega, a-1) i_R(\omega, a-1)\Big) m(\omega, a-1) = \\ \Big(z(\omega, a) + z_S\,(\omega, a) + x\Big) m(\omega, a).\end{aligned} \tag{A.9}$$

This equation states that inflows into state (ω, a) must equal outflows from that same state. Inflows are given by incumbent innovations on their own product, either generated themselves or acquired from a related startup. Any innovation triggers an outflow.

Moreover, the distribution must hold

$$\sum_{a=1}^{+\infty} m(\omega, a) = \tau(\omega). \tag{A.10}$$

Equations (A.9) and (A.10) define a system that can easily be solved for the distribution m once we know all innovation rates and implementation and noticing probabilities.

A.3 Aggregate variables and the growth rate

Aggregate variables Product market clearing implies $Y_t = C_t + I_t^{\text{Inc}} + I_t^{\text{Startup}} + S_t$, where

$$I_t^{\text{Inc}} = Y_t \Bigg[\xi_I^{\text{ext}} x^{\psi} + \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) \Bigg(\xi_I^{\text{int}} \big(z(\omega, a)\big)^{\psi} + \dots \\ \cdots + x_S \Big[(1-\gamma) s_R(\omega, a) \kappa_I \big(i_R(\omega, a)\big)^{\theta} + \gamma s_U \kappa_I \big(i_U(\omega)\big)^{\theta} \Big] \Bigg) \Bigg], \tag{A.11}$$

$$I_t^{\text{Startup}} = Y_t \left[x_S \left(\xi_S + \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) \Big(\gamma(1 - s_U) + (1-\gamma)(1 - s_R(\omega, a)) \Big) \kappa_S \big(i_S(\omega)\big)^{\theta} \right) \right], \tag{A.12}$$

$$S_t = Y_t \left[\chi(s_U)^{\varphi} + \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) \chi \big(s_R(\omega, a)\big)^{\varphi} \right], \tag{A.13}$$

are the aggregate costs of incumbent innovation, startup innovation, and search, respectively, all of which scale linearly in Y_t. This shows that consumption is always proportional to Y_t, and thus grows at the same rate. Note that government tax revenues are included in consumption, as they are rebated lump-sum to the representative consumer.

Growth rate By equation (29), output, consumption and wages grow at the rate of Q_t. In BGP, this growth rate is constant. Using (27), note we can write:

$$Q_t = \exp\left(\int_0^1 \omega_j \ln(\omega_j) \, \mathrm{d}j \right) \underbrace{\exp\left(\int_0^1 \omega_j \ln(q_{jt}) \, \mathrm{d}j \right)}_{\equiv \mathcal{Q}_t}. \tag{A.14}$$

As the first term is constant, we have $g = \frac{\dot{Q}_t}{Q_t} = \frac{\dot{\mathcal{Q}}_t}{\mathcal{Q}_t}$. For a small period of time dt, we have

$$\ln \mathcal{Q}_{t+dt} = \int_0^1 \omega_j \ln q_{j,t+dt} \mathrm{d}j. \tag{A.15}$$

We can partition this integral by spending shares and by the technology gap of each product j depending on whether or not the product in innovated on from time t to $t+dt$, so that

$$\begin{aligned}\ln \mathcal{Q}_{t+dt} &= \sum_{\omega\in\Omega}\left[\sum_{a=1}^{+\infty}\left(\int_{j\in\mathcal{J}_t(\omega,a)} \omega \ln\left(\lambda q_{j,t}\right)\mathrm{d}j + \int_{j\notin\mathcal{J}_t(\omega,a)} \omega \ln\left(q_{j,t}\right)\mathrm{d}j\right)\right] \\ &= \ln \mathcal{Q}_t + \sum_{a=1}^{+\infty}\left(\int_{j\in\mathcal{J}_t(\omega,a)} \omega \ln(\lambda)\mathrm{d}j\right),\end{aligned}$$

where $\mathcal{J}_t(\omega,a) \subset [0,1]$ is the set of products in state (ω,a) that are innovated on from periods t to $t+dt$. In BGP, a fraction $\Big(x + z(\omega,a) + z_S(\omega,a)\Big)dt$ of products with spending share ω and technology gap a see their productivity increase by a factor λ. The remaining fraction of products keeps its productivity unchanged. Therefore,

$$\frac{\ln(\mathcal{Q}_{t+dt}/\mathcal{Q}_t)}{dt} = \ln(\lambda)\sum_{\omega\in\Omega}\omega\left[\sum_{a=1}^{+\infty} m(\omega,a)\Big(x + z(\omega,a) + z_S(\omega,a)\Big)\right]. \tag{A.16}$$

Taking the limit as $dt \to 0$ and rearranging terms gives equation (30).

Decomposition formula Finally, we derive the growth decomposition formula. We can write the growth rate in (30) as

$$g = \ln(\lambda)\Big(\mathcal{I} + x_S\mathcal{P}\Big), \tag{A.17}$$

where $\mathcal{I}$ and $\mathcal{P}$ are the sales-weighted averages of innovations of incumbent ideas and of the implementation probability of startup ideas, respectively, that is:

$$\mathcal{I} \equiv x + \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m(\omega,a)\omega z(\omega,a) \tag{A.18}$$

$$\begin{aligned}\mathcal{P} \equiv \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m(\omega,a)\omega\Big(&(1-\gamma)\, s_R(\omega,a) i_R(\omega,a) + \gamma s_U i_U(\omega) + \dots \\ &\cdots + \big(1-(1-\gamma)\, s_R(\omega,a) - \gamma s_U\big) i_S(\omega)\Big)\end{aligned} \tag{A.19}$$

To get the decomposition of the change in the growth rate between any two BGPs

labeled A and B, we can simply rewrite (A.17) as:

$$\frac{g^B}{g^A} = \frac{\ln(\lambda)\mathcal{I}^A}{g^A}\frac{\mathcal{I}^B}{\mathcal{I}^A} + \frac{\ln(\lambda)x_S^A\mathcal{P}^A}{g^A}\frac{x_S^B}{x_S^A}\frac{\mathcal{P}^B}{\mathcal{P}^A}, \tag{A.20}$$

which is equation (32) with $\text{Share}_{\text{inc}}^A \equiv \frac{\ln(\lambda)\mathcal{I}^A}{g^A}$.

A.4 Additional tables and figures

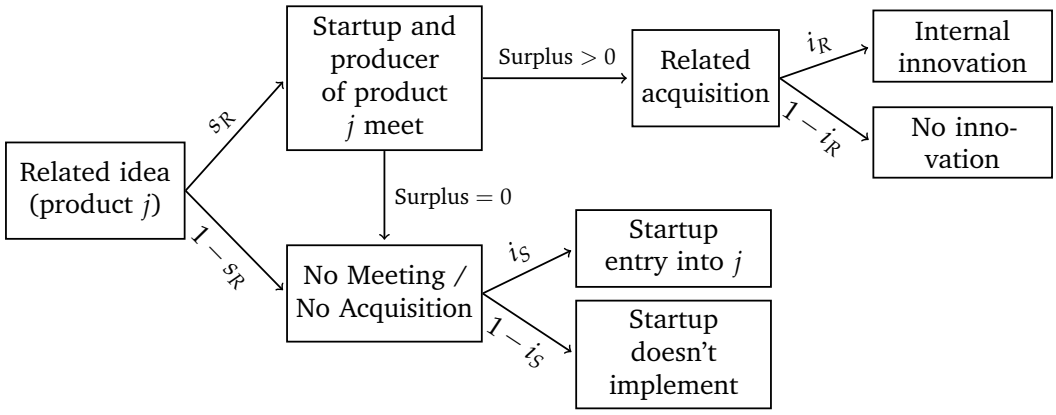

Figure A.1: Timing of events for a *related* startup idea (on product j) within a period $(t, t+dt)$.

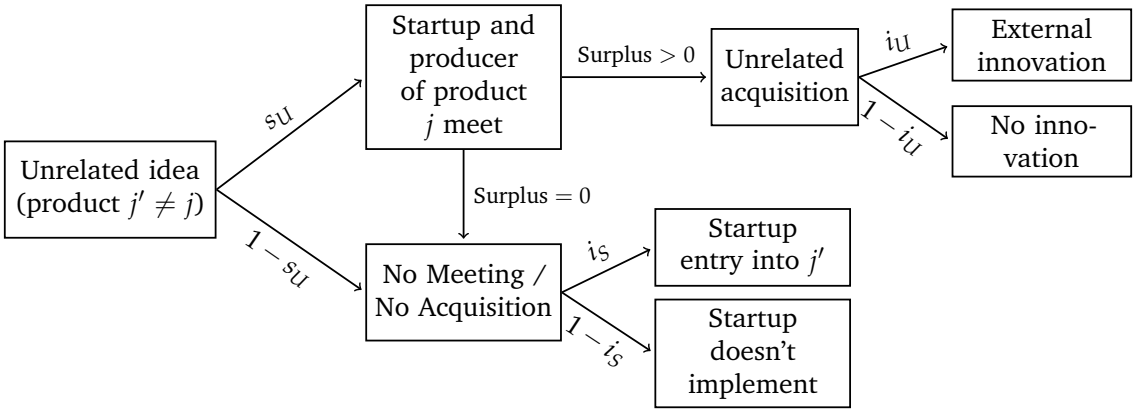

Figure A.2: Timing of events for an *unrelated* startup idea (on product $j' \neq j$) within a period $(t, t+dt)$.

B Data appendix

B.1 Matching startups to patents

Our data sources are fully described in the main text. This section contains further details on the way in which we match startup acquisition targets from SDC Platinum to their patents in PatentsView.

Our matching procedure is based on firm names (the only common information across the two datasets). To perform the match, we first standardize firm names, by removing commonly used endings such as "CO", "CORP" or "CORPORATION". This yields some exact matches, but their number is relatively small. Therefore, we complement this by a manual matching exercise, considering all names of startup acquisition targets in SDC Platinum and checking whether there is a closely corresponding firm name in PatentsView. This catches many cases with minor variations in the spelling of firm names, such as for instance "FRONTIER PHARMA" and "FRONTIER PHARMACY".

B.2 Matching acquired to control patents

This section provides further details on how we match treated to control patents. Our sample of treated patents has around 600 entries, as shown in the main text. For each of these treated patents, we aim to find a set of control patents that looks at 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 four categories: patent application year, patent technological subcategory, patent inventors' location (U.S. state), and target firm founding year. In other words, we limit the pool of potential control patents to the ones that have the same patent application year, the same patent technological subcategory, the same inventor location within the U.S. and where the firm was founded the same year as the treated patent.[41]

Second matching stage Conditional on satisfying the criteria of the first matching stage, we estimate a propensity score matching algorithm based on the number of inventors that took part in developing the patent. At the expense of smaller sample sizes, in robustness tests

[41] Data on the founding year of patenting firms comes from the "Startup Patenting" project of Ewens and Marx, available online at https://foundingpatents.com/.

we also have accounted for additional features like the number of claims[42] or text-based patent features extracted by Arts *et al.* (2021).

We perform the propensity score matching allowing for up to 10 control patents for each treated patent. In particular, if the set of control patents that have the closest propensity score to the treatment patent is smaller than 10 patents, then we keep all of them. If this set of closest control patents exceeds 10 patents, then we keep a random sample of 10 units. For this reason, some treated patents might have 5 or 6 control patents, while others have the full 10 control patents. In robustness tests, we allow for a lower or higher number of control patents, but the main result is unaffected.

B.3 Summary statistics

Table B.1 compares treated patents to control patents after our propensity score matching. A total of 479 patents (out of 596) can be matched to control patents. By construction, these treated patents have similar mean values for the number of inventors than the 3225 control patents. Furthermore, the t-test for the difference in means is also not significant for variables that were not used in the matching, such as the number of claims, citations received after 1 year, the number of new bigrams in the patent text and the degree of textual novelty (the latter two variables are from Arts *et al.* (2021)). Overall, this table supports the claim that both sets of patents are similar across a large number of observables.

Table B.1: Comparison of Matching Patents

	Control Patents		Treatment Patents			
	Mean	Obs	Mean	Obs	Difference	p-value
Number of Inventors	3225	2.64	479	2.64	0.01	0.57
Number of Claims	3225	22.13	479	21.37	-0.76	0.28
Citations Received after 1 Year	3225	0.22	479	0.23	0.02	0.68
New Bigrams	3008	4.68	435	4.13	-0.54	0.34
Textual Novelty	3008	0.97	435	0.97	-0.00	0.28

Notes: This table compares treated to up to 10 control patents for the variables incorporated in the propensity score matching exercise. All variables are defined in Section B.2.

[42] Claims specify 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).

B.4 Empirical results: Robustness exercises

Table B.2 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 do not receive any boost in citations after acquisition compared to control patents of the same matched patent pair. Column (2) increases the time window to 7 years prior to the acquisition and 7 years after the acquisition (our baseline window considered 5 years). Column (3) adds non-granted patents to the dependent variable. Column (4) includes industry-year fixed effects instead of year fixed effects. Column (5) uses an OLS regression where the dependent variable is in levels. Finally, column (6) uses an OLS regression where the dependent variable is in logs. In all specifications, we again conclude that there is no differential impact when looking at the full sample.

Finally, Table B.3 analyzes whether there is a differential effect for the pharmaceutical industry. The results of Cunningham *et al.* (2021), documenting killer acquisitions in this industry, could suggest a more negative effect of acquisitions on citations here. To test this, we consider a triple difference model like the one in equation (35), where we replace the *Related* dummy with a dummy equal to 1 if the acquisition target belongs to the pharmaceutical industry. We do find evidence for a negative effect, although our results are not statistically significant.

Table B.2: Robustness: main regression results

	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.721***	0.831***	0.721***	0.785***	0.206*	0.0952***
	(0.139)	(0.159)	(0.139)	(0.137)	(0.109)	(0.0296)
Acquired	-0.209	-0.217	-0.213	-0.0323	0.0544	0.0142
	(0.437)	(0.444)	(0.437)	(0.414)	(0.123)	(0.0428)
Acquired x Post	0.262	0.408	0.265	0.0904	0.000653	0.0168
	(0.525)	(0.546)	(0.525)	(0.500)	(0.315)	(0.0975)
Matched Patents FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Baseline	✓					
Longer Trend		✓				
Adds non-granted Patents			✓			
Industry-Year FE				✓		
OLS levels					✓	
OLS logs						✓
Observations	22,119	26,243	22,119	32,547	33,064	33,064

Notes: We use a Poisson estimator. The dependent variable is the number of citations received at the patent-year level. *Acquired* is a dummy taking value 1 for treated patents, and *Post* is a dummy taking value 1 for post-acquisition years. For ease of comparison, column (1) repeats our baseline result. Column (2) expands the window of years to 7 years prior acquisition until 7 years post acquisition. Column (3) adds non-granted patents to the dependent variable. Column (4) includes industry-year fixed effects instead of year fixed effects. Column (5) is an OLS regression in levels. Finally, column (6) 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.3: Robustness: the pharmaceutical industry

	(1)	(2)	(3)
Post	0.721***	0.715***	0.716***
	(0.139)	(0.143)	(0.143)
Acquired	-0.209	-0.242	-0.242
	(0.437)	(0.485)	(0.485)
Acquired x Post	0.262	0.325	0.324
	(0.525)	(0.575)	(0.575)
Post x D(Pharma)		0.0965	0.0941
		(0.377)	(0.376)
Acquired x D(Pharma)		0.370	0.376
		(0.803)	(0.804)
Acquired x Post x D(Pharma)		-0.951	-0.945
		(1.114)	(1.117)
Matched Patents FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	22,119	22,119	22,119

Notes: We use a Poisson estimator. The dependent variable is the number of citations received at the patent-year level. *Acquired* is a dummy taking value 1 for treated patents, *Post* is a dummy taking value 1 for post-acquisition years, and *Pharma* is a dummy taking value 1 if an acquisition target belongs to the pharmaceutical industry. For ease of comparison, column (1) repeats our baseline result. In Column (2), the *Pharma* dummy is defined with respect to its 3-digit NAICS code (325), and in Column (3), with respect to its 4-digit NAICS code (3254). Standard errors are clustered at the target firm level. * significant at 10%; ** significant at 5%; *** significant at 1%.

C Numerical and calibration details

C.1 BGP solution algorithm

To solve for our model's BGP equilibrium, we use the following algorithm:

1. Guess a value for the startup rate, $x_S > 0$.
2. Given this guess, solve for the value function of incumbents through a value function iteration on the HJB equation (8):
   (a) Guess a value function, $v : \Omega \times \mathbb{A} \to \mathbb{R}$, where $\mathbb{A} \equiv \{1, 2, \ldots, a_{max}\}$ and a_{max} is a sufficiently large integer.[43]
   (b) Using equations (9), (10), (12), (14), (17), (19) and (22), deduce from this guess the policy functions z, x, i_S, i_U, s_U, i_R, and s_R, as well as the acquisition surpluses b_U and b_R, using (15) and (20).
3. Using the innovation rates obtained in step 2 and the guess for the startup rate x_S, compute the joint distribution of products across spending shares and technology gaps, $m(\omega, a)$, as described in Appendix A.2.
4. Compute the value of creating a startup v_S (the right-hand side of equation (25)). If
$$\left| \frac{\xi_S - v_S}{v_S} \right| < 10^{-4},$$
the equilibrium has been found. Otherwise, update the x_S guess 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, using the equations in Section 2.2.6. This block structure greatly simplifies the solution of the model.

C.2 Computing moments in the model

This section shows how we compute each of the model moments targeted in our calibration. We discuss moments in the same order as they appear in Section 4.

[43] To make sure this bound is not binding, in practice we check that $\sum_{\omega \in \Omega} m(\omega, a_{max}) \approx 0$.

Share of unrelated startup acquisitions First, we compute the percentage of all startups that are acquired, which holds

$$S^{\text{Acq}} \equiv \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) \Big(\gamma s_U + (1-\gamma) s_R(\omega, a) \Big). \tag{C.1}$$

Then, our first moment is

$$M_1 \equiv \frac{\gamma s_U}{S^{\text{Acq}}}. \tag{C.2}$$

Acquisition frequency The acquisition frequency is the ratio between the arrival rate of startup acquisitions and the mass of incumbent firms. Therefore, it holds

$$M_2 = \frac{x_S S^{\text{Acq}}}{1/\overline{n}}, \tag{C.3}$$

where $\overline{n}$ is the average number of products per firm. As the mass of products is equal to 1, the mass of firms is the inverse of the average number of products per firm. To compute this statistic, we rely on a simulation of a large cohort of firms, as described in Appendix C.3.

Effect of acquisitions on the implementation probability of startup ideas In our model, each startup idea is characterized by the characteristics of the product that it applies to. Moreover, we know the implementation probability for each idea, depending on whether it is acquired or not. Therefore, we can compute the causal effect of acquisition on (the natural logarithm of) the implementation probability by taking the difference between the log implementation probabilities of incumbents and startups for each product type, and weighting these by the share of startup acquisitions represented by each product type (which governs how often a certain type of acquisition would show up in an empirical sample). It is convenient to compute these statistics first separately for related and unrelated acquisitions. We get

$$\beta_R \equiv \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m_R(\omega, a) \Big[\ln \big(i_R(\omega, a) \big) - \ln \big(i_S(\omega) \big) \Big], \tag{C.4}$$

$$\beta_U \equiv \sum_{\omega \in \Omega} \tau(\omega) \Big[\ln \big(i_U(\omega) \big) - \ln \big(i_S(\omega) \big) \Big] = \frac{1}{\theta - 1} \ln \left(\frac{\kappa_S}{\kappa_I} \right), \tag{C.5}$$

respectively, where

$$m_R(\omega, a) \equiv \frac{m(\omega, a) s_R(\omega, a)}{\sum_{\omega'} \sum_{a'} m(\omega', a') s_R(\omega', a')} \tag{C.6}$$

denotes the distribution of acquired related startups over spending shares and technology gaps.[44] The overall causal effect is a weighted average of these two statistics, with the weight given by the share of unrelated acquisitions (our first moment).

$$M_3 = M_1 \cdot \beta_U + (1 - M_1) \cdot \beta_R. \tag{C.7}$$

Average acquisition premium For each acquisition, the acquisition premium is the ratio of the acquisition price to the outside option of the startup. The average premium can be computed as a weighted average of the premia for related and unrelated acquisitions, so that

$$M_4 = M_1 \cdot AcqPrem_U + (1 - M_1) \cdot AcqPrem_R, \tag{C.8}$$

where

$$AcqPrem_R \equiv (1-\alpha) \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m_R(\omega, a) \max \left[\frac{b_R(\omega, a)}{\kappa_S(\theta - 1)(i_S(\omega))^\theta}, 0 \right], \tag{C.9}$$

$$AcqPrem_U \equiv (1-\alpha) \sum_{\omega \in \Omega} \tau(\omega) \max \left[\frac{b_U(\omega)}{\kappa_S(\theta - 1)(i_S(\omega))^\theta}, 0 \right], \tag{C.10}$$

Relationship between acquisition probability and firm size Using our panel of simulated firms (see Appendix C.3), we compute the yearly acquisition probability for each simulated firm f and year t as

$$AcqProb_{ft} \equiv x_S \left(n_{ft} \gamma s_U + (1 - \gamma) \sum_{j=1}^{n_{ft}} s_R(\omega_{jt}, a_{jt}) \right) \tag{C.11}$$

where n_{ft} is the number of products of the firm at the beginning of year t. Then, we compute the relationship between firm size and the acquisition probability (the counterpart of regression (33) in the data) by running the following OLS regression on our simulated data:

$$AcqProb_{ft} = \alpha + \beta_{\text{Acq}} \ln \left(Sales_{ft} \right) + u_{ft} \tag{C.12}$$

where $Sales_{ft} = \sum_{i=1}^{n_{ft}} \omega_{ft}$ are the firm's (normalized) sales, and $M_5 = \beta_{\text{Acq}}$.

Average implementation probability of non-acquired startups The probability of startup implementation (i.e. entry) conditional on no acquisition is computed as

$$M_6 = \frac{1}{1 - S^{\text{Acq}}} \left[\sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) \Big(\gamma \big(1 - s_U\big) + (1 - \gamma) \big(1 - s_R(\omega, a)\big) \Big) i_S(\omega) \right]. \tag{C.13}$$

[44] For unrelated acquisitions, the distribution over spending shares is $\tau(\omega)$, as search for them is undirected.

This is the ratio of the arrival rate of startup ideas successfully implemented by startups, and the arrival rate of non-acquired startups. S^{Acq}, the share of acquired startups, was defined in equation (C.1).

Growth and sources of growth The aggregate growth rate g is our seventh moment M_7. To decompose this growth rate, we divide the aggregate innovation rate in three components:

$$\begin{aligned} I_{\text{Inc}}^{\text{CD}} &= x + x_S \gamma s_U \sum_{\omega \in \Omega} \tau(\omega) \omega i_U(\omega), \\ I_{\text{Inc}}^{\text{OwnProd}} &= \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) \omega \left[z(\omega, a) + x_S \Big((1-\gamma) s_R(\omega, a) i_R(\omega, a) \Big) \right], \\ I_{\text{Ent}} &= x_S \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} \left[m(\omega, a) \omega \Big(\gamma (1 - s_U) + (1-\gamma)\big(1 - s_R(\omega, a)\big) \Big) i_S(\omega) \right], \end{aligned}$$

where $g = \ln(\gamma)(I_{\text{Inc}}^{\text{CD}} + I_{\text{Inc}}^{\text{OwnProd}} + I_{\text{Ent}})$. Thus, the share of growth accounted for by entrants is

$$M_8 = \frac{I_{\text{Ent}}}{I_{\text{Inc}}^{\text{CD}} + I_{\text{Inc}}^{\text{OwnProd}} + I_{\text{Ent}}} \tag{C.14}$$

and the share of creative destruction in incumbent's contribution to growth is

$$M_9 = \frac{I_{\text{Inc}}^{\text{CD}}}{I_{\text{Inc}}^{\text{CD}} + I_{\text{Inc}}^{\text{OwnProd}}}. \tag{C.15}$$

Exit rate In the BGP of our model, the exit rate of innovative firms is equal to the entry rate, given by

$$M_{10} = \frac{x_S}{1/\bar{n}} \left[\sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) i_S(\omega) \Big(\gamma (1 - s_U) + (1-\gamma)(1 - s_R(\omega, a)) \Big) \right]. \tag{C.16}$$

This is the ratio of the arrival rate of entrants to the mass of incumbent firms.

C.3 Simulation

To compute some of the moments described in Appendix C.2, we need to know the distribution of products (and their characteristics) across firms. For this, we simulate a cohort of 5,000 entering firms over a period of $T = 200$ years. For each firm, we track its age and its product portfolio, namely the number of products and the spending share ω

and technology gap a of each product. This yields the joint distribution of firm product portfolios and firm age. Integrating over age, we obtain the unconditional distribution of product portfolios across firms, which we can use for instance to compute the average number of products per firm.[45]

Precisely, we assume that all 5,000 firms start off as entrants, holding one product with technology gap $a = 1$. The share of entrants with a low spending share product is set to match the equivalent number in the BGP equilibrium. Then, we update the product portfolio of each firm by using the product-level transition rates implied by the policy functions of the BGP equilibrium.[46]

The simulation yields various moments related to the joint distribution of firms and products, such as the distribution of firm sales (see Figure 5), the firm age distribution (see Figure C.1, left) and the distribution of products per firm (see Figure C.1, right).

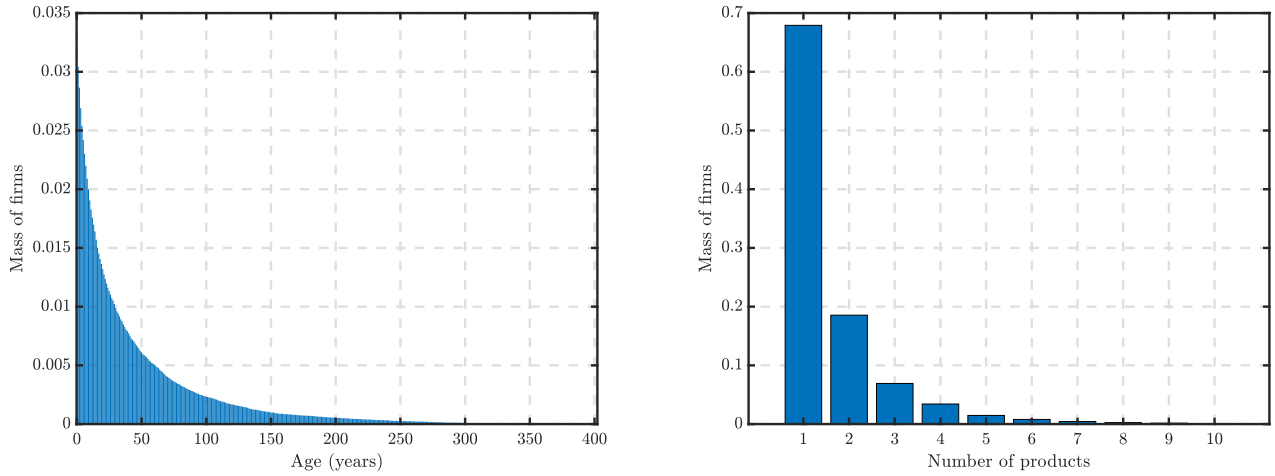

Figure C.1: The age distribution and the distribution of the average number of products per firm in the simulation of the model

C.4 Global identification

This section lays out a global identification test for the calibration in Section 4. We use a method first developed by Daruich (2022).

Denote our vector of parameters by $\boldsymbol{\theta} \in \boldsymbol{\Theta} \subset \mathbb{R}_{+}^{10}$, where $\boldsymbol{\Theta}$ is the parameter space. First, we create a large 10-dimensional hyper-cube $\mathcal{P} \subset \boldsymbol{\Theta}$ 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

[45] This simulation approach follows Garcia-Macia *et al.* (2019).
[46] In doing so, we partition time into short periods of length $dt = 1/150$.

uniform partitions of the space. For a sufficiently large number of Sobol draws, this routine efficiently and comprehensively covers $\mathcal{P}$. 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. After N Sobol draws (in practice, $N \approx 13$ million), we have a $N \times 10$ matrix $\boldsymbol{R}$ of results and a $N \times 10$ matrix $\boldsymbol{P} \in \mathcal{P}$ of the corresponding parameters.

The model-generated data contained in the $(\boldsymbol{R}, \boldsymbol{P})$ matrices can then be exploited to obtain information about identification. First, for each parameter $p \in \boldsymbol{\theta}$, 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 quasi-random variation in all the remaining 9 parameters. Using this fact, we 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 can 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.

Figure C.2 presents the results from this procedure, where we have associated each targeted moment with the parameter that the moment most plausibly identifies (the same pairing as in Table 5 and in Section 4). All in all, we find that all the parameters of the model are very well-identified by all four criteria.

C.5 Quantitative results: Robustness exercises

In Section 5.3, we discussed the implications of targeting different values for the acquisition premium, the causal impact of acquisitions on patent citations, and the contribution of entrants to growth. In this section, we provide the same analysis for all other targeted moments. Figures C.3 and C.9 show the results for these exercises, using the same format as in the main text. They indicate that overall, our results are robust to reasonable variations in these moments.

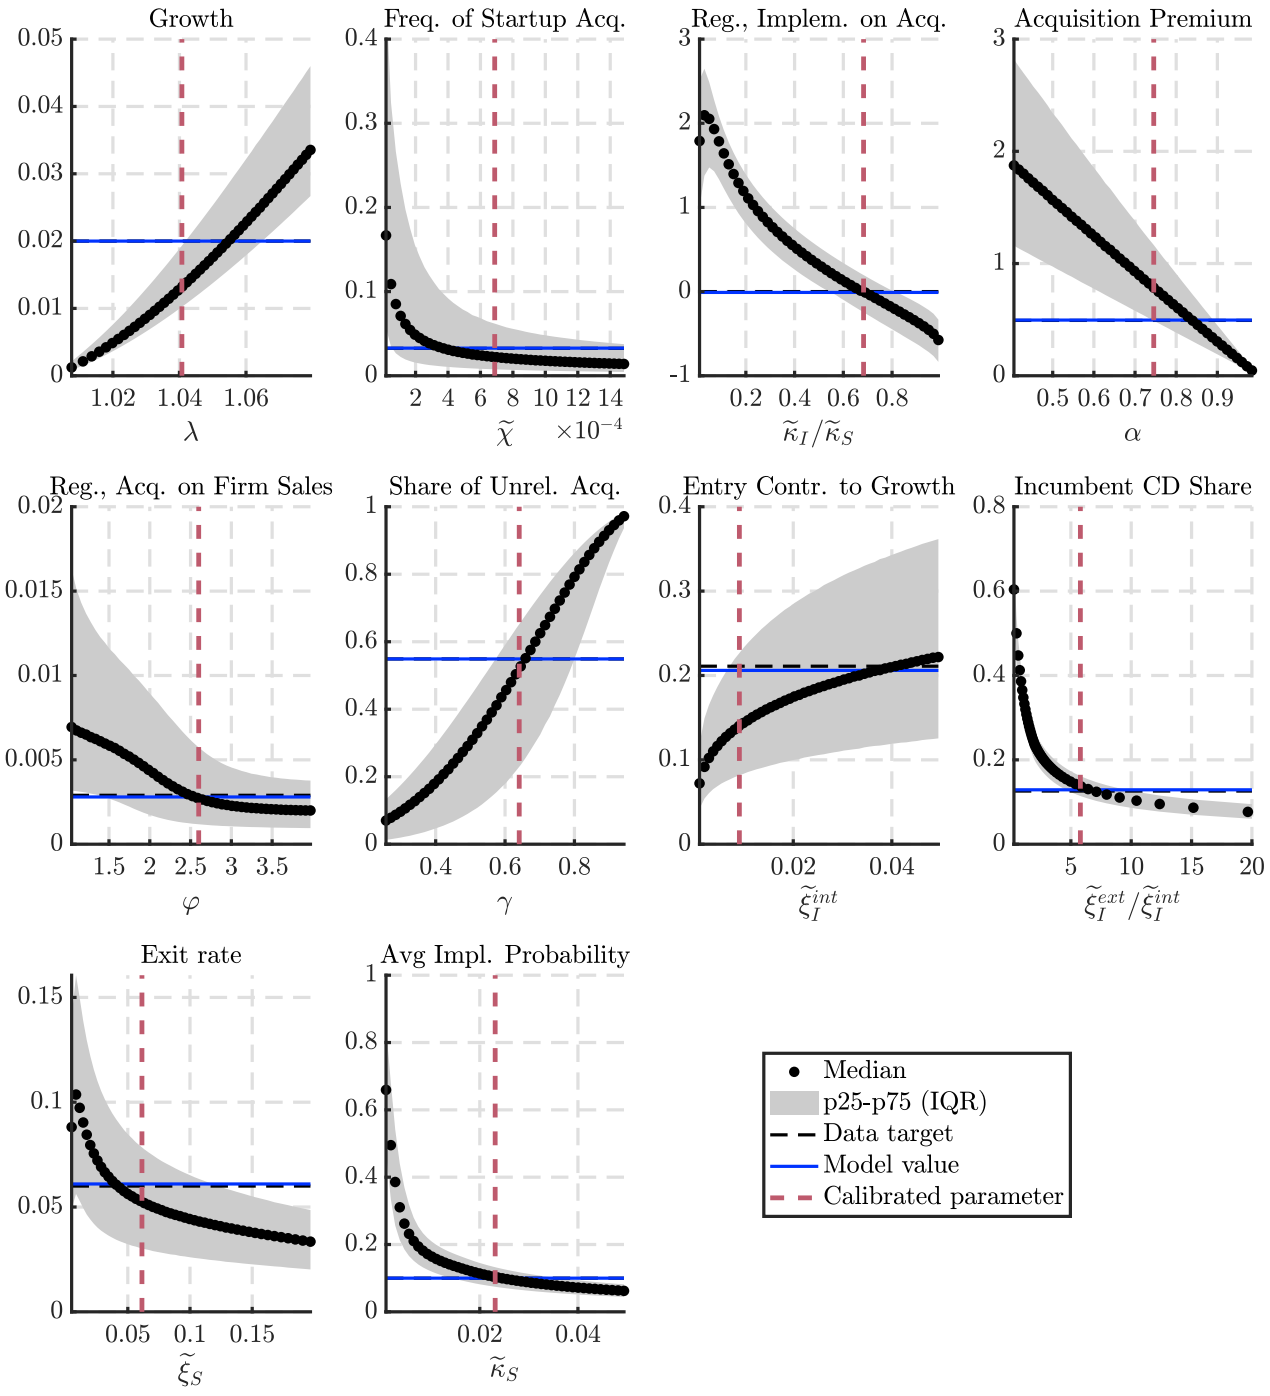

Figure C.2: Global identification results based on 13 million model solutions.
Notes: For each parameter-moment pairing, the black dots are the median of the distribution generated by random variation in all the other parameters, and the gray area covers the inter-quartile range of the distribution. The dashed horizontal line marks the empirical target, and the vertical line marks the calibrated parameter value. The cost scale parameters χ, ζ_I^{int}, ζ_I^{ext}, κ_S and κ_I reported in the main text have been rescaled relative to the ones reported in this figure to match the costs of achieving a 10% success probability, that is $\widetilde{\chi} = \chi \cdot 0.1^{\varphi}$, $\widetilde{\zeta}_I^{\text{int}} = \zeta_I^{\text{int}} \cdot 0.1^{\psi}$, $\widetilde{\zeta}_I^{\text{ext}} = \zeta_I^{\text{ext}} \cdot 0.1^{\psi}$, $\widetilde{\kappa}_S = \kappa_S \cdot 0.1^{\theta}$, and $\widetilde{\kappa}_I = \kappa_I \cdot 0.1^{\theta}$.

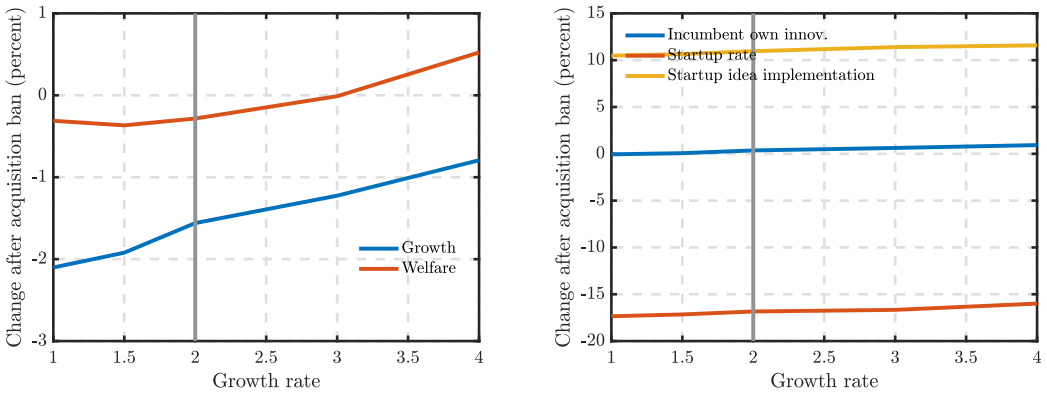

Figure C.3: Robustness to different targets for the aggregate growth rate. *Notes:* See Figure 7.

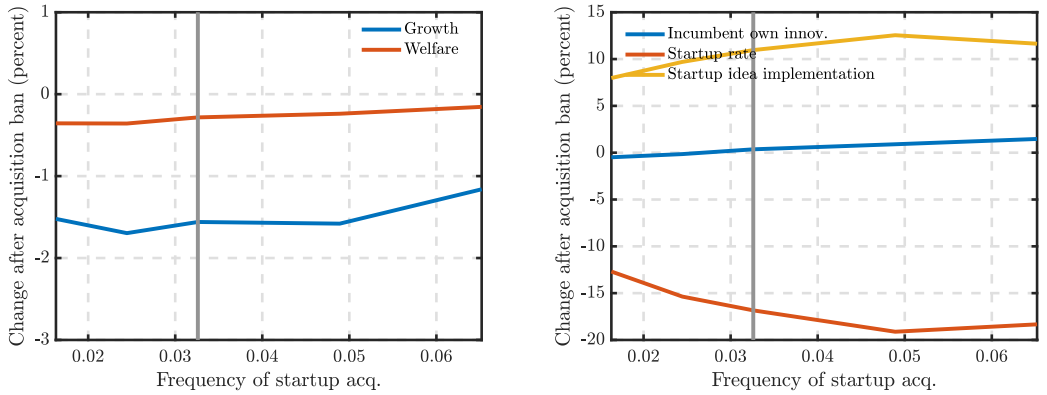

Figure C.4: Robustness to different targets for the startup acquisition frequency. *Notes:* See Figure 7.

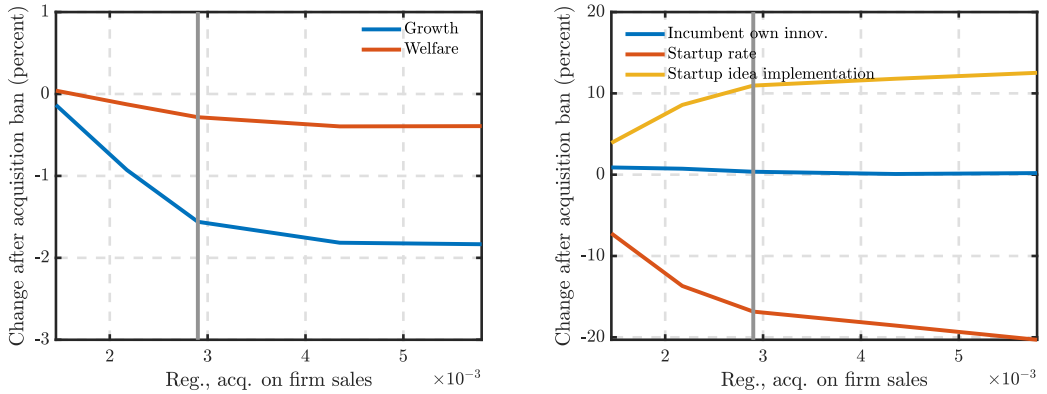

Figure C.5: Robustness to different targets for the link between acquisition and acquirer size. *Notes:* See Figure 7.

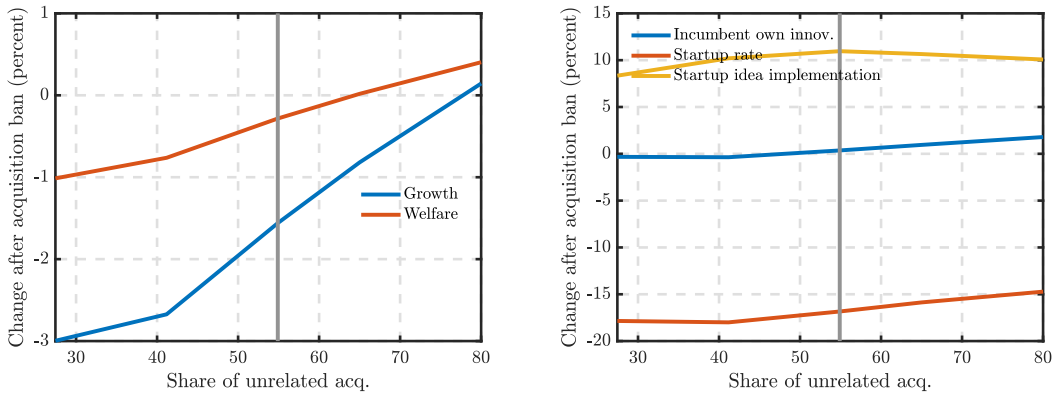

Figure C.6: Robustness to different targets for the share of unrelated acquisitions. *Notes:* See Figure 7.

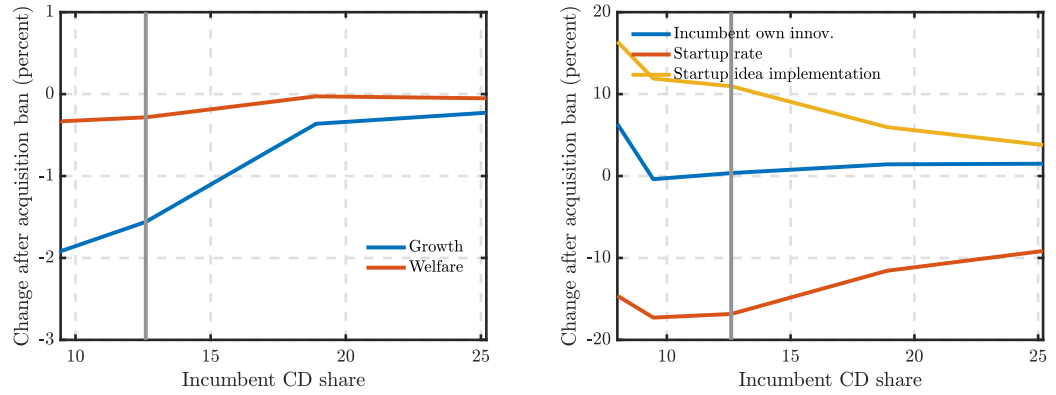

Figure C.7: Robustness to different targets for the share of creative destruction in incumbents' own innovations. *Notes:* See Figure 7.

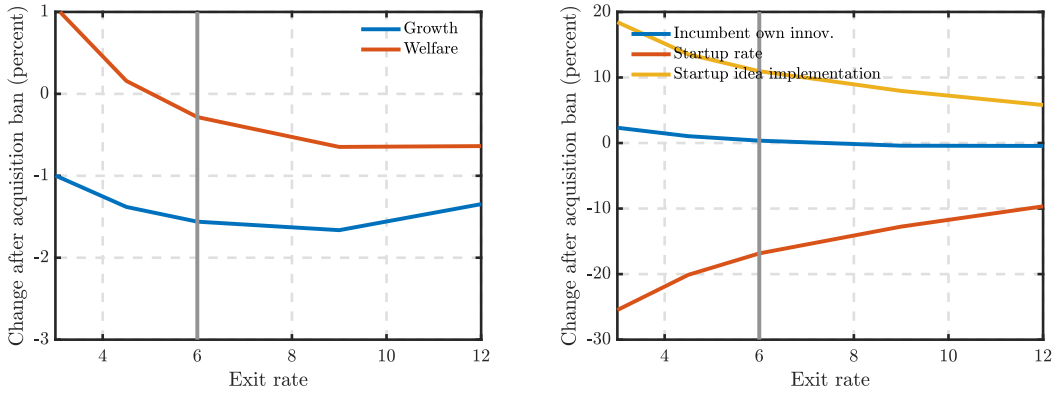

Figure C.8: Robustness to different targets for the exit rate. *Notes:* See Figure 7.

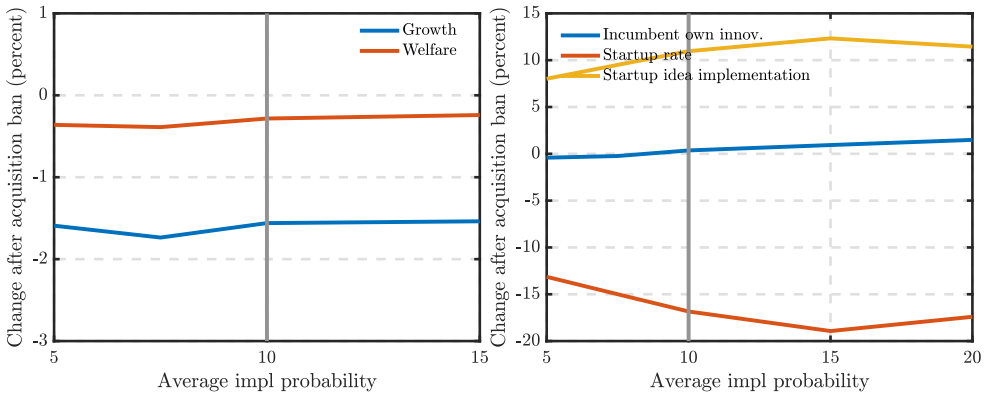

Figure C.9: Robustness to different targets for the average implementation probability of startups. *Notes:* See Figure 7.

In particular, Figure C.9 illustrates the claim made in Section 4: our results are virtually unchanged irrespective of whether the target for the implementation probability is 5, 10 or 15%. This irrelevance is due the fact that the calibrated level of i_S does not matter for policy. In all equilibrium conditions, outcomes do not depend on i_S, but on the product $x_S \cdot i_S$, which is effectively the product-level exit rate (abstracting from acquisitions, which are infrequent events). The target for the implementation rate of startups pins down the relative contributions of i_S and x_S, but this decomposition is irrelevant for growth outcomes. With heterogeneous products and acquisitions, these arguments no longer hold exactly. However, as Figure C.9 shows, the irrelevance of the target for the implementation probability of startups still continues to apply quantitatively.