The Effects of Startup Acquisitions on Innovation and Economic Growth*

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

February 9, 2026

Abstract

Innovative startups are frequently acquired by large incumbents. 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. This 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 taxes on startup acquisitions lower the startup rate, but increase incumbent innovation as well as the implementation rate of startup ideas. Banning killer acquisitions, a policy that appears desirable in partial equilibrium, yields virtually no welfare gains in general equilibrium. The optimal policy instead imposes high taxes on startup acquisitions (reducing their frequency by more than half) and raises consumption-equivalent welfare by 0.48%.

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 four anonymous referees, 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, Joseba Martinez, Enrico Moretti, Jane Olmstead-Rumsey, Michael Peters, Joshua Weiss, Liangjie Wu, and participants at several institutions and conferences for helpful comments. Pau Roldan-Blanco acknowledges financial support from the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2024-000915-S), funded by MCIN/AEI/10.13039/501100011033, and from Ayudas para Proyectos de Investigación en Ciencias Sociales 2024, funded by Fundación Ramón Areces. Part of this research was conducted while Pau Roldan-Blanco was an Associated Faculty member at CEMFI. This research utilised Queen Mary's Apocrita HPC facility, supported by QMUL Research-IT (http://doi.org/10.5281/zenodo.438045).

†University of California in Merced, CEPR, and Max Planck Institute for Innovation and Competition in Munich; 5200 North Lake Rd. Merced, CA 95343, United States, cfons-rosen@ucmerced.edu.

‡Universitat Autònoma de Barcelona, Barcelona School of Economics, and CEPR; Department of Economics and Economic History, UAB Campus Building B, 08193 Bellaterra, Barcelona, Spain, pau.roldan@uab.cat.

§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.[1] However, many successful startups never grow into large independent firms, as they are acquired early on by older incumbents. For incumbents, in turn, startup acquisitions are a routine activity: several of the largest U.S. firms have bought hundreds of startups over the last decade.[2]

In recent years, regulators have viewed these operations with increasing skepticism. In 2020, the Federal Trade Commission (FTC) announced an inquiry into high-profile startup acquisitions and filed lawsuits against two large incumbents, Alphabet and Meta.[3] 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*".[4]

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), eliminating innovative ideas that threaten their business. Other analysts have pushed back on this view, pointing out that acquisitions may motivate founders to create startups in the first place and that incumbents might be 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. Our paper aims to provide the first such comprehensive assessment.

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] Decker *et al.* (2014, 2016) show that startups account for a large share of job creation and are overrepresented among high-growth firms. Sedláček and Sterk (2017), Sterk *et al.* (2021) and De Haas *et al.* (2022), among others, emphasize that startup growth is largely driven by ex-ante heterogeneity.

[2] According to the Federal Trade Commission (FTC), the Tech giants Alphabet, Amazon, Apple, Meta and Microsoft acquired more than 600 small firms between 2010 and 2019 (FTC, 2021).

[3] A report on this inquiry was released in September 2021 (FTC, 2021). The FTC sued the companies in October and December 2020, asking Meta to undo its acquisitions of Instagram and WhatsApp.

[4] See https://itif.org/publications/2022/01/31/platform-competition-and-opportunity-act-solution-search-problem. The act never came to a vote (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 both "related" startups (i.e., startups with an idea applying to one of the incumbent's products) and "unrelated" startups (i.e., startups with an idea applying to 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 after an acquisition, an incumbent may invest more or less into implementation than the startup that came up with the idea.

To regulate startup acquisitions, we assume that a government can impose proportional taxes on acquisition prices, with different rates for related and unrelated acquisitions. These taxes affect the economy's long-run growth rate, the most important driver of consumer welfare, through three margins: (i) changes in the startup rate (i.e., the number of startups created in a given year); (ii) changes in the own innovation rate of incumbents; and (iii) changes in the implementation rate of startup ideas. Taxes also have a welfare-relevant effect on the consumption share of output.[5]

The policy debate, with its focus on killer acquisitions, has largely concentrated on just one of these margins: the link between acquisitions and the implementation rate of startup ideas. However, a full assessment of startup acquisitions needs to go beyond this direct effect, and take into account general equilibrium spillovers on other margins. This is the goal of our paper. To achieve it, we rely on quantitative predictions from our model, which we discipline by calibrating its parameters 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 from Compustat. Focusing on the period 2000-2020, we define startup acquisitions as deals in which the target is younger than six years, following Guzman and Stern (2020). With this data, we generate several calibration targets, including an estimate of the causal effect of acquisitions on the implementation of startup ideas. In our baseline analysis, we proxy idea implementation by patent citations. That is, 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,

[5] Finally, taxes affect the level of output, as they shift the distribution of markups and thereby affect the efficiency of the labor allocation across firms. However, this effect turns out to be negligible in our calibration.

we match each acquired patent to a group of similar control patents. We find that the average acquisition has no significant effect on citations. However, we find a negative effect on citations if the startup and its acquirer are related, i.e., belong to the same NAICS 4-digit industry. This is consistent with our model, where related acquisitions are subject to a replacement effect (Arrow, 1962): incumbents have less incentives to implement related ideas than startups, as this displaces some of their pre-existing profits. Our results are robust to many different specifications, including alternative proxies for idea implementation.

We calibrate the model to match our regression evidence and other relevant moments, such as the average acquisition premium or the frequency of startup acquisitions.[6] We then study the effect of acquisition taxes, computing the transition path of the economy from a no-policy balanced growth path (BGP) to a BGP with acquisition taxes.

We first study the most salient policy in the public debate: a ban on related acquisitions, implemented through an infinite tax. A partial equilibrium assessment of this policy (where we assume that related acquisitions disappear, but all other firm decisions are unaffected), suggests a 4 basis point increase in long-run growth and a 1.29% increase in consumption-equivalent welfare. However, general equilibrium effects eliminate this gain. As acquisition prices always exceed a startup's outside option of remaining independent, the acquisition ban reduces incentives for startup creation, and leads to a 16.5% lower startup rate. For incumbents, there are positive and negative effects. Related 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 the ban also reduces creative destruction, which increases incumbent value and innovation incentives. In our calibration, the net effect is a small increase in incumbents' own innovation. However, this does not fully compensate for the fall in the startup rate, and long-run growth falls by 1 basis point. As the consumption share of output slightly rises (with fewer resources being used for startup creation and search), the overall welfare effect of the ban is virtually zero.

Other policies can achieve better results. We find that welfare is maximized for a high 215% tax on related acquisitions and a lower 30% tax on unrelated acquisitions. This policy cuts the frequency of acquisitions by more than half, and increases consumption-equivalent welfare by 0.48%. Qualitatively, these taxes have the same effects as a related

[6] To match the findings from our regressions, our model needs to assume that incumbents have lower idea implementation costs than startups. Then, unrelated acquisitions increase the implementation probability of startup ideas, canceling out the negative effect of related acquisitions. While we take differences in implementation costs to be exogenous, they could be micro-founded by economies of scale and scope, a larger customer base, better access to capital markets, or greater business experience. For example, Ignaszak and Sedláček (2023) model how customer accumulation depends on firm characteristics.

acquisition ban: a decrease in the startup rate, an increase in incumbent own innovation, and an increase in startup idea implementation. However, the policy maximizes the positive response of incumbent own innovation (the largest driver of aggregate growth), so that overall growth now increases by 1 basis point, driving the welfare gain. The key to understanding this result is to note that the policy is biased in favor of high-markup incumbents. These firms still find acquisitions profitable despite taxes, and even see an increase in their share of the acquisition surplus. Thus, the tax creates incentives for all other incumbents to increase their markups through innovation.

To conclude our analysis, we show how results depend on calibration targets. We find that acquisition taxes are more desirable when acquisitions are frequent, when they are likely to be killer acquisitions, and when acquisition premia are high. Overall, through a wide range of alternative values for our calibration targets, the optimal policy always involves a high tax on related startup acquisitions, and welfare gains around 0.5%.

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 R&D. Phillips and Zhdanov (2013) instead argue that acquisitions stimulate innovation by small firms aiming to be acquired. Using data on listed 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 several 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 (Federico *et al.*, 2017; Cabral, 2018; Bourreau *et al.*, 2018; Bryan and Hovenkamp, 2020; Fumagalli *et al.*, 2020; Kamepalli *et al.*, 2020; Denicolò and Polo, 2021; Brutti and Rojas, 2022; Callander and Matouschek, 2022; Eisfeld, 2023; Letina *et al.*, 2024). 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 growth.[7] 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. Cavenaile *et al.* (2021), Guthmann and Rahman (2025) 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 welfare effects of startup acquisitions, taking into account new general equilibrium spillovers on the startup rate and on the implementation of ideas.[8] 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 Section 2 presents our model. Section 3 describes our micro data, lays out stylized facts, and estimates the effects of acquisitions on the implementation of startup ideas. Section 4 presents our calibration. Section 5 highlights the channels through which acquisitions affect welfare, and discusses our quantitative results. Section 6 concludes.

2 Model

Our model describes the links between acquisitions and innovation. We build on heterogeneous-firm Schumpeterian growth models, but introduce two new elements: startup acquisitions and a distinction between invention and 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

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

where $\rho > 0$ is the 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

[7] Pellegrino (2025) and Cao and Zhu (2025) 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).

[8] Weiss (2022), Pearce and Wu (2022) and Wei *et al.* (2023) also develop endogenous growth models with acquisitions or patent trade, but these do not feature all the margins studied in our paper. Celik *et al.* (2022) study the effects of information frictions in the merger market on innovation and business dynamism.

supplies $L > 0$ units of labor at the wage w_t. She 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

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

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, taking values in a finite set Ω and holding $\sum_{\omega\in\Omega} \omega\phi(\omega) = 1$, where $\phi(\omega) \in [0, 1]$ denotes the (exogenous) mass of products with spending share ω.

Each product can potentially be produced by a large number of multi-product firms. A 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 paragraphs describe their innovation technologies.

Incumbent innovation Incumbent firms innovate 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 $\zeta_I^{\text{ext}} n x^{\psi} Y_t$ units of the final good, an incumbent that currently produces n products generates a Poisson arrival rate nx for an external innovation. An external innovation applies to a randomly chosen product in the interval $[0,1]$ and enables the incumbent to improve its leading productivity by a factor λ. In equilibrium, the innovating incumbent then overtakes production of this product, displacing the previous producer through creative destruction.

Startup innovation 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 in the interval $[0,1]$. 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.[9] 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 randomly chosen product $j' \in [0,1]$.[10]

The probability that potential acquirers notice the startup is endogenous, and depends on their effort in monitoring the startup scene.[11] 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 is 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, search costs are increasing and convex, and scale with aggregate output to ensure balanced growth.

Conditional on a meeting, the incumbent may acquire the startup, by transferring p_{jt}^A

[9] 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.

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

[11] 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.

units of the final good (henceforth, the acquisition price) to the startup in exchange for the startup exiting forever and handing over its idea to the incumbent. The incumbent also needs to pay an acquisition tax which is proportional to the price, at a rate τ_R for a related acquisition and τ_U for an unrelated acquisition.[12] The acquisition price is determined through Nash bargaining, where the incumbent has a bargaining weight $\alpha \in (0, 1)$.

After an acquisition, the incumbent 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).[13]

2.2 Balanced Growth Path Equilibrium

2.2.1 Household decisions, markups and profits

In this section, we solve for our model's balanced growth path (BGP) equilibrium, in which all aggregate variables grow at a constant rate $g > 0$. We assume throughout that parameter values are such that the equilibrium features positive entry.

On the BGP, consumption growth holds the Euler equation $\frac{\dot{C}_t}{C_t} \equiv g = r - \rho$. The demand for any product j is given by $y_{jt} = \omega_j \frac{Y_t}{p_{jt}}$, where p_{jt} denotes the price of the product. Bertrand competition implies that for any product j, only the firm with the highest productivity produces, charging a limit price which is equal to the marginal cost of its follower (the firm with the second-highest productivity). Followers are former incumbents, displaced when the current incumbent overtook production. Denoting by q^F_{jt} 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 $\lambda^{a_{jt}} \equiv q_{jt}/q^F_{jt}$. Then, we have

$$p_{jt} = \lambda^{a_{jt}} \frac{w_t}{q_{jt}}, \tag{4}$$

where $\lambda^{a_{jt}}$ is the markup. A product with spending share ω_j and technology gap a_{jt} generates a profit

$$\pi_t\left(\omega_j, a_{jt}\right) = \omega_j \left(1 - \lambda^{-a_{jt}}\right) Y_t. \tag{5}$$

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}.

[12] The proceeds of these taxes are rebated lump-sum to the representative consumer.

[13] Figure A.1 in the Online Appendix summarizes the timing of events for startup ideas.

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, implementation probabilities for acquired startup ideas, as well as an external innovation rate x.

External innovations scale linearly in firm size (as in Klette and Kortum, 2004), which implies that an incumbent's value function is additively separable across products. Moreover, as profits and costs are proportional to aggregate output, the value function is also proportional to aggregate output. Thus, as shown in Online Appendix A.2, 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

$$V_t(\mathbf{n}, n) = \sum_{j=1}^{n} v(\omega_j, a_j) Y_t \tag{6}$$

for some time-invariant product-level value function v. The Hamilton-Jacobi-Bellman (HJB) equation for the product-level value function is

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

$$\underbrace{+ x \sum_{\omega' \in \Omega} \phi(\omega') v(\omega', 1)}_{\text{External innovation}} - \underbrace{\tilde{x} v(\omega, a)}_{\substack{\text{Creative destruction}\\ \text{by external innovation}}} + \underbrace{x_S v_R(s_R, \omega, a)}_{\substack{\text{Startup has idea on}\\ \text{incumbent's product}}} + \underbrace{x_S \gamma v_U(s_U)}_{\substack{\text{Unrelated startup}\\ \text{matched to incumbent}}} \Bigg\}. \tag{7}$$

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

The HJB equation shows how the discounted value of a product of type (ω, a) changes over time. In every instant, the incumbent collects profits and spends on innovation costs and startup search. At Poisson rate z, it makes an internal innovation, increasing its technology gap by one step. At rate x, it makes an external innovation, allowing it to overtake a new product with spending share ω' (chosen at random from the distribution of spending shares), and a technology gap 1. At rate $\tilde{x}$, the incumbent 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 denoted by $v_R(s_R, \omega, a)$, which we derive below. Finally, at rate $x_S\gamma$, the incumbent is matched to an unrelated startup. This changes its value by $v_U(s_U)$, which we again derive below.

From problem (7), we find that the product-level internal innovation rate holds

$$z(\omega, a) = \left(\frac{v(\omega, a+1) - v(\omega, a)}{\xi_I^{\text{int}} \psi} \right)^{\frac{1}{\psi - 1}}. \tag{8}$$

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

$$x = \left(\frac{\sum_{\omega' \in \Omega} \phi(\omega') v(\omega', 1)}{\xi_I^{\text{ext}} \psi} \right)^{\frac{1}{\psi - 1}}. \tag{9}$$

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.

2.2.3 Unrelated acquisitions

First, consider the case in which an unrelated startup is matched to the incumbent. This startup has an idea on a product j', with spending share ω' and technology gap a'. Conditional on the match, the incumbent either 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, the incumbent might acquire the startup, thereby increasing its value.

To determine whether an acquisition occurs, and at what price, we solve the Nash bargaining problem. 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, it is the value of the problem

$$\max_{i_S} \left\{ i_S v(\omega', 1) - \kappa_S i_S^{\theta} \right\}, \tag{10}$$

where i_S is the startup's optimal implementation probability in the absence of a meeting. When the 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 an implementation probability equal to

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

and its outside option can be rewritten as $\kappa_S (\theta - 1) (i_S(\omega'))^{\theta}$.

When an acquisition occurs, the continuation value of the startup is the acquisition price,

$p_U^A(\omega')$. For the incumbent, conditional on an acquisition price $p_U^A(\omega')$, the continuation value is the value of the problem

$$\max_{i_U} \left\{ i_U v(\omega', 1) - \kappa_I i_U^\theta \right\} - (1+\tau_U)\, p_U^A(\omega'). \tag{12}$$

The incumbent therefore implements the acquired idea with probability

$$i_U(\omega') = \left(\frac{v(\omega', 1)}{\kappa_I \theta} \right)^{\frac{1}{\theta - 1}}. \tag{13}$$

As was the case for the startup, this implementation probability only depends on the product's spending share ω'. Putting all of these elements together, we can now write the Nash bargaining problem between the startup and the incumbent as

$$\begin{aligned} \max_{p_U^A(\omega')} \quad & \left(\kappa_I\, (\theta - 1) \left(i_U(\omega') \right)^\theta - (1+\tau_U)\, p_U^A(\omega') \right)^\alpha \left(p_U^A(\omega') - \kappa_S\, (\theta - 1) \left(i_S(\omega') \right)^\theta \right)^{1-\alpha} \\ \text{such that} \quad & p_U^A(\omega') \geq \kappa_S\, (\theta - 1) \left(i_S(\omega') \right)^\theta . \end{aligned} \tag{14}$$

The constraint ensures that the startup is paid at least its outside option. Problem (14) has a solution (i.e., an acquisition will take place) if and only if there is a surplus, this is, if the joint value of both firms after the acquisition exceeds the sum of their outside options. In Online Appendix A.2, we show that the surplus $b_U(\omega')$ holds

$$b_U(\omega') = (1 + \alpha \tau_U)\, (\theta - 1) \left[\frac{\kappa_I}{1+\tau_U} \left(i_U(\omega') \right)^\theta - \kappa_S \left(i_S(\omega') \right)^\theta \right]. \tag{15}$$

It is easy to show that the surplus is non-negative if and only if $\tau_U \leq \left(\frac{\kappa_S}{\kappa_I} \right)^{\frac{1}{\theta-1}} - 1$. In particular, without a tax, an unrelated acquisition occurs if the incumbent has lower implementation costs than the startup (i.e., $\kappa_I \leq \kappa_S$). In this case, the idea is more valuable in the hands of the incumbent, and an acquisition allows this transfer to take place.

When the surplus is non-negative, the acquisition price is

$$p_U^A(\omega') = \kappa_S\, (\theta - 1) \left(i_S(\omega') \right)^\theta + (1 - \widetilde{\alpha}_U) b_U(\omega'), \quad \text{with } \widetilde{\alpha}_U \equiv \frac{\alpha(1+\tau_U)}{1+\alpha\tau_U}. \tag{16}$$

The acquisition price is the sum of the startup's outside option and the startup's effective share of the surplus. This share, equal to $1 - \widetilde{\alpha}_U$, is decreasing in both the incumbent's bargaining power α and the tax rate τ_U.

The incumbent receives the remaining share $\widetilde{\alpha}_U$ of the surplus. As it notices unrelated startups with probability s_U, and the ideas of these startups apply to a randomly chosen product, the expected change in the incumbent's value from being matched to an unrelated startup is

$$v_U(s_U) = s_U\, \widetilde{\alpha}_U \sum_{\omega' \in \Omega} \phi(\omega') \max\Big(0, b_U(\omega')\Big). \tag{17}$$

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

$$s_U = \left[\frac{\widetilde{\alpha}_U \gamma x_S}{\chi \varphi} \left(\sum_{\omega' \in \Omega} \phi(\omega') \max\left(0, b_U(\omega')\right)\right)\right]^{\frac{1}{\varphi - 1}}. \tag{18}$$

Search effort increases in the surplus, and does not depend on incumbent characteristics.

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 startup is not matched to the incumbent, the latter faces potential displacement either from the startup itself, with probability $(1 - s_U) i_S(\omega)$, or from an unrelated incumbent acquiring the startup, with probability $s_U i_U(\omega)$.[14] In both cases, implementation makes the related incumbent lose its entire value. Hence, incumbent value decreases in expectation by $[s_U i_U(\omega) + (1 - s_U) i_S(\omega)] v(\omega, a)$.

When the related startup is matched to the incumbent, but goes unnoticed (with probability $1 - s_R$), the incumbent is displaced with probability $i_S(\omega)$. However, if the incumbent notices the startup, it might acquire it. To characterize this decision, we solve another bargaining problem. The startup's outside option is still defined by problem (10). The incumbent's outside option is $-i_S(\omega) v(\omega, a)$, its expected value loss in case it does not acquire. Finally, if the incumbent acquires the startup, its continuation value is

$$\max_{i_R} \left\{ i_R \Big(v(\omega, a+1) - v(\omega, a) \Big) - \kappa_I i_R^{\theta} \right\} - (1 + \tau_R) p_R^A(\omega, a), \tag{19}$$

where $p_R^A(\omega, a)$ is the related acquisition price. Thus, the incumbent implements the idea of a related startup with probability

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

[14] 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.

Given this, the Nash bargaining problem is

$$\max_{p_R^A(\omega,a)} \left(\kappa_I\left(\theta-1\right)\left(i_R(\omega)\right)^{\theta}-\left(1+\tau_R\right)p_R^A(\omega,a)+i_S(\omega)v(\omega,a)\right)^{\alpha}\left(p_R^A(\omega,a)-\kappa_S\left(\theta-1\right)\left(i_S(\omega)\right)^{\theta}\right)^{1-\alpha}$$
$$\text{such that} \quad p_R^A(\omega,a)\geq\kappa_S\left(\theta-1\right)\left(i_S(\omega)\right)^{\theta}. \tag{21}$$

A related acquisition takes place if and only if it generates a surplus. As we show in Online Appendix A.2, the surplus $b_R(\omega,a)$ holds

$$b_R(\omega,a)=\left(1+\alpha\tau_R\right)\left[\left(\theta-1\right)\left(\frac{\kappa_I}{1+\tau_R}\left(i_R(\omega)\right)^{\theta}-\kappa_S\left(i_S(\omega)\right)^{\theta}\right)+\frac{i_S(\omega)v(\omega,a)}{1+\tau_R}\right]. \tag{22}$$

The first summand in the square brackets captures the surplus generated by lower implementation costs of incumbents, as for unrelated acquisitions. However, there is now a second summand, and a second source of surplus: avoiding displacement preserves incumbent rents. Indeed, displacement requires costly implementation and resets the technology gap to 1, while the status quo maintains the current (potentially higher) technology gap for free.

When the surplus is non-negative, the acquisition price is

$$p_R^A(\omega,a)=\kappa_S\left(\theta-1\right)\left(i_S(\omega)\right)^{\theta}+(1-\widetilde{\alpha}_R)b_R(\omega,a), \qquad \text{with} \quad \widetilde{\alpha}_R\equiv\frac{\alpha(1+\tau_R)}{1+\alpha\tau_R}. \tag{23}$$

This is the sum of the startup's outside option and its effective surplus share $1-\widetilde{\alpha}_R$, which decreases in both the incumbent's bargaining power α as well as the tax τ_R.

As the incumbent receives the remaining share $\widetilde{\alpha}_R$ of the surplus, the expected change in its value when a startup has an idea on its product can finally be written as:

$$v_R(s_R,\omega,a)=-\Big(\gamma s_U i_U(\omega)+\left(1-\gamma s_U\right)i_S(\omega)\Big)v(\omega,a)+\left(1-\gamma\right)s_R\widetilde{\alpha}_R\max\Big(0,b_R(\omega,a)\Big). \tag{24}$$

The first summand is the change in value without an acquisition, i.e., the product of the current value and the displacement probability (a weighted average of the implementation probability of startups and unrelated incumbents). The second summand is the product of the probability of noticing the startup, $(1-\gamma)s_R$, and the surplus from an acquisition. From this, the search effort for related startups holds

$$s_R(\omega,a)=\left(\frac{\widetilde{\alpha}_R\left(1-\gamma\right)x_S\max\left(0,b_R(\omega,a)\right)}{\chi\varphi}\right)^{\frac{1}{\varphi-1}}. \tag{25}$$

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

We can now wrap up the dynamic problem of incumbents. Replacing all our results into equation (7), product value holds

$$v(\omega,a) = \frac{\pi(\omega,a) + (\psi-1)\left(\xi_I^{\text{int}} z(\omega,a)^{\psi} + \xi_I^{\text{ext}} x^{\psi}\right) + \chi\,(\varphi-1)\left(s_R(\omega,a)^{\varphi} + s_U^{\varphi}\right)}{\rho + x + x_S\Big(\gamma s_U i_U(\omega) + (1-\gamma s_U) i_S(\omega)\Big)}, \tag{26}$$

where we used the fact that, in equilibrium, $\widetilde{x} = x$ (all incumbents choose the same external innovation rate). Product 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.

2.2.5 Implementation of startup ideas

Before proceeding, it is useful to compare implementation choices of startups and incumbents. From our optimality conditions, we have

$$\left(\frac{i_R(\omega,a)}{i_S(\omega)}\right)^{\theta-1} = \underbrace{\frac{v(\omega,a+1)-v(\omega,a)}{v(\omega,1)}}_{\text{Replacement effect}} \quad \underbrace{\frac{\kappa_S}{\kappa_I}}_{\text{Cost differences}} \quad \text{and} \quad \left(\frac{i_U(\omega)}{i_S(\omega)}\right)^{\theta-1} = \frac{\kappa_S}{\kappa_I}. \tag{27}$$

We will show later that the value function is concave in the technology gap. Thus, for related acquisitions, there is a replacement effect: innovation implies a smaller increase in value for an incumbent than for a startup.[15] All else equal, this implies that related acquisitions lower the implementation probability of startup ideas. However, there is also a cost effect: if incumbents have sufficiently low implementation costs relative to startups, related acquisitions may still increase the implementation probability. For unrelated acquisitions, in turn, differences in implementation probabilities are exclusively driven by cost differences.

2.2.6 Startup creation

In equilibrium, the cost of creating a startup is equal to the expected benefit:

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

[15]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.

$$\ldots + \underbrace{(1-\gamma)\, s_R(\omega,a)\,(1-\widetilde{\alpha}_R)\, b_R\,(\omega,a) + \gamma s_U\,(1-\widetilde{\alpha}_U)\, b_U(\omega)}_{\text{Expected surplus from related and unrelated acquisitions}}\Bigg]. \tag{28}$$

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 (28) takes an expectation using the joint distribution of products over these two states, denoted by $m(\omega, a)$. This distribution is an endogenous outcome, driven by innovation and acquisition choices, and we characterize it in Online Appendix A.3. The startup always obtains its outside option, and if it is noticed by an incumbent, it also obtains a share of the acquisition surplus.[16] Thus, all else equal, acquisitions increase the value of creating a startup.

2.2.7 Aggregate outcomes

To close the model, we solve for aggregate outcomes. Combining the demand function and the product price, we get that the incumbent producing product j demands $\omega_j \lambda^{-a_{jt}} \frac{Y_t}{w_t}$ units of labor. Using labor market clearing, we obtain the aggregate labor share:

$$\frac{w_t L}{Y_t} = \frac{1}{\mathcal{M}}, \qquad \text{where } \mathcal{M} \equiv \left(\int_0^1 \omega_j \, \lambda^{-a_{jt}} dj \right)^{-1}. \tag{29}$$

The labor share is the inverse of the aggregate markup, $\mathcal{M}$, defined as a harmonic weighted average of product-level markups. Aggregate output holds

$$Y_t = \frac{\mathcal{M}}{\mathcal{W}} Q_t L, \;\; \text{where} \;\; Q_t \equiv \exp\left(\int_0^1 \omega_j \ln(\omega_j q_{jt}) dj \right), \; \mathcal{W} \equiv \exp\left(\int_0^1 \omega_j \ln(\lambda^{a_{jt}}) dj \right). \tag{30}$$

The term $\mathcal{M}/\mathcal{W} < 1$ is an efficiency wedge, as heterogeneous markups trigger a misallocation of labor across products. Without misallocation, aggregate output would simply be equal to the product of aggregate productivity Q_t and labor supply L.

Equation (30) shows that output grows at the rate of aggregate productivity, g. Then, as we show in Online Appendix A.4,

$$g = \ln(\lambda)\left(x + \sum_{\omega\in\Omega} \sum_{a=1}^{+\infty} \left[m(\omega,a)\omega\Big(z(\omega,a) + x_S i(\omega,a) \Big)\right]\right), \quad \text{with} \tag{31}$$

$$i(\omega,a) \equiv \; (1-\gamma)\, s_R(\omega,a) i_R(\omega,a) + \gamma s_U i_U(\omega) + \big(1-(1-\gamma)\, s_R(\omega,a) - \gamma s_U\big) i_S(\omega).$$

Growth is proportional to the aggregate innovation rate. The innovation rate is, in turn,

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

the sum of the own innovation rates of incumbents (x and z) and the innovation rate from startup ideas (the product of the startup rate x_S and the average implementation probability of startup ideas i). Rates are weighted by spending shares ω, which in equilibrium are equal to product sales shares. Finally, product market clearing implies that output is used for consumption (C_t), innovation by incumbents and startups (I_t), and search (S_t), so that $Y_t = C_t + I_t + S_t$. Online Appendix A.4 shows that the shares C_t/Y_t, I_t/Y_t and S_t/Y_t are all constant on the BGP.

This completes the characterization of our model's BGP equilibrium. Online Appendix A.5 explains 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.[17] Figure 1 plots the product-level value function v and the internal innovation arrival rate z. We assume throughout 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 internal innovation rate, in turn, depends on the increments of the value function. Therefore, it is increasing in ω and decreasing in a.

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

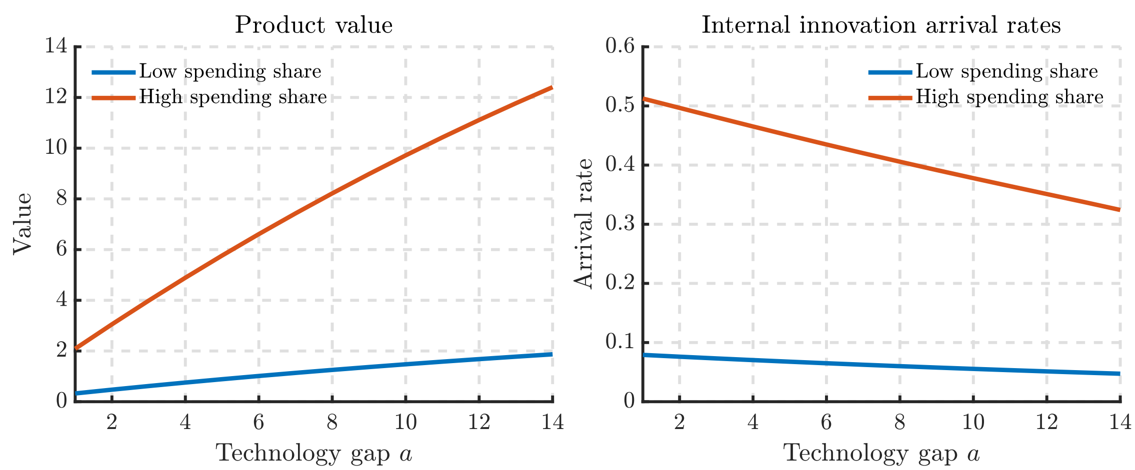

The left panel of Figure 2 plots the noticing probabilities s_R and s_U as a function of product characteristics. For unrelated startups, the noticing probability is constant. For related startups, noticing probabilities are increasing in the spending share ω and the

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

technology gap a. Indeed, as discussed before, related acquisitions have two benefits. First, when incumbents have lower implementation costs, 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. Both benefits are increasing in ω and a, explaining the behavior of noticing probabilities.

Figure 2: Noticing and implementation probabilities of startup ideas by product type.

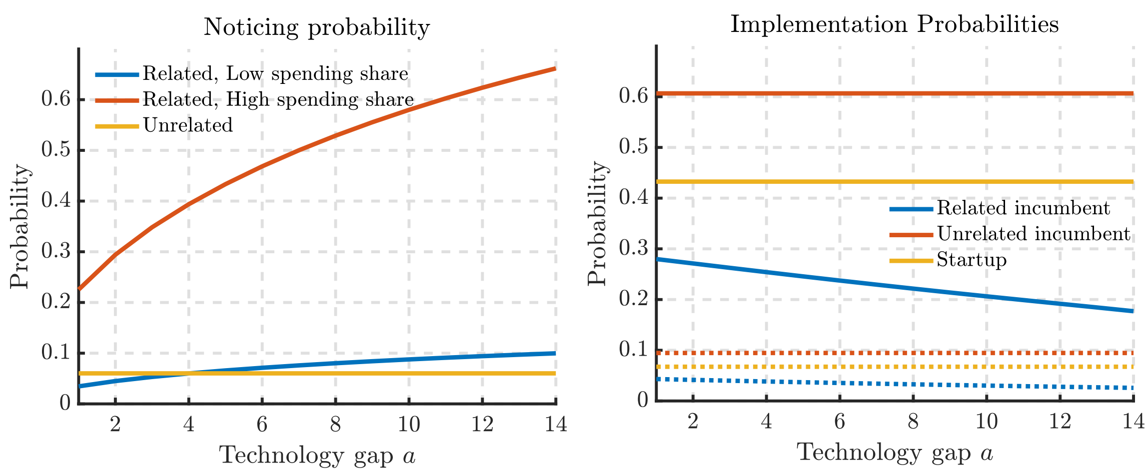

Notes: In the right panel, full lines refer to high-spending-share products, while dotted lines refer to low-spending-share products.

The right panel of Figure 2 plots implementation probabilities for startup ideas, depending on whether the startup itself, a related or an unrelated incumbent invest in implementation. 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 overall lower implementation incentives, and all related acquisitions are killer acquisitions (that is, every related acquisition lowers the probability of the acquired idea being implemented). Moreover, as the marginal benefit of implementation decreases in the technology gap, related acquirers with high technology gaps are least likely to implement startup ideas.

The previous discussion shows that many incumbent decisions 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 innovation and acquisition choices of all firms: for instance, Figure 3 shows that the average technology gap is higher for high-spending-share products, as they receive more innovation.

This completes the description of our model. In order to use it as a framework for a quantitative analysis, we now need to map it to the data. To do so, the next section introduces a new micro-level data set, which allows us to generate calibration targets.

Figure 3: The equilibrium distribution of technology gaps

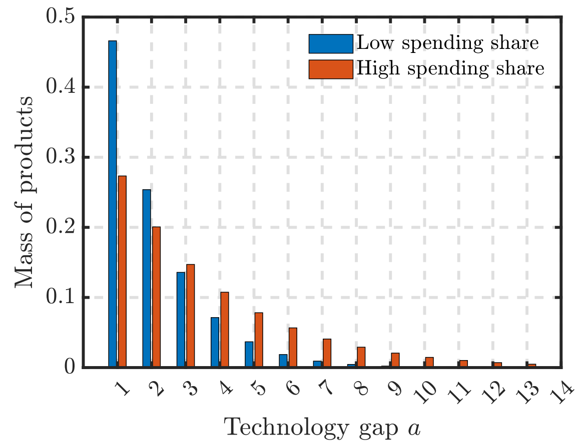

Notes: This figure plots the distribution of technology gaps a across products, within each spending share category. Online Appendix A.3 describes how we solve for this distribution.

3 Empirical analysis

3.1 Data

Data sources To track startup acquisitions, we rely on SDC Platinum, provided by the London Stock Exchange Group. 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 acquiring 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.[18]

In order to measure the innovation activity of startups, we rely on patent data from the USPTO's PatentsView database. PatentsView contains the universe of US patents since 1976. It provides information on the firm owning each patent (name and assignee code), as well as forward and backward citations, the number of claims assigned to the patent, the identity of the inventors, and patent technology classes.

Finally, to recover some balance sheet characteristics of listed acquirers, we use Compustat. Online Appendix B contains further information on all data sources.

[18]Unfortunately, SDC does not always report the foundation date of acquisition targets. We are able recover some of these missing values by relying on the work of Ewens and Marx (2024), who provide founding dates for many patenting firms. Online Appendix B.1 contains further details.

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.* (2025), 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 first standardize firm names using the `stnd_compname` routine, developed by Wasi and Flaaen (2015) for Stata. Then, we match standardized firm names in SDC to standardized firm names in PatentsView, using exact matches only (see Online Appendix B.2 for details).

3.2 Descriptive statistics

We observe around 200 startup acquisitions per year in our sample. There is no strong time trend: after a low point in 2001 (with just over 100 deals), acquisitions steadily increased until 2014, reaching over 300 deals. After 2014, there has been a slowdown.[19]

Table 1 summarizes some 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 $35 million, but the distribution is highly skewed, with the mean exceeding the median by a factor of 8. Some acquired startups (around 10% of those with a deal value) were already publicly listed before their acquisition. For these, SDC provides an 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 50.7%.[20]

For our empirical analysis, we define related acquisitions as acquisitions in which the target and the acquirer share the same NAICS 4-digit industry code. As Table 1 shows, this is the case for 39% of startup acquisitions in our sample. Finally, we find that 620 target firms hold patents before being acquired (with a median number of two). Measured by citations, these patents are of higher quality than others: they received on average 7.3 citations during their first five years, as opposed to 3.1 citations for all other patents.

We will use several of these stylized facts to calibrate our model. However, we first

[19]Figure B.1 in the Online Appendix plots the number of startup acquisitions over time. In an earlier version of this paper, we used data from Guzman and Stern (2020), which indicated that about 4% of all patenting startups are acquired within their first six years. This percentage has no clear time trend either.

[20]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%.

turn to the central part of our empirical analysis, assessing the effect of acquisitions on the implementation of startup ideas.

Table 1: Summary statistics for the startup acquisition sample

	N	Mean	Median	p10	p90
Acquisition year	4615	2011	2012	2002	2019
Target age (years)	4615	3.6	3.8	1.5	5.6
Deal value (millions)	1266	287	35	2	522
Acquisition Premium (percent)	110	50.7	33.3	-4.4	130.2
Same industry	4615	0.39	0.00	0.00	1.00
Acquirer sales (millions)	1653	13867.02	1176.58	45.46	37905.00
Acquirer employment (thousands)	1653	33.09	4.00	0.19	75.05
Number of Patents	620	3.43	2.00	1.00	7.50
Citations (Acquired Patents)	2127	7.27	2.00	0.00	19.00
Citations (Non-Acquired Patents)	3697278	3.08	0.00	0.00	7.00

Notes: This table considers a sample of deals between 2000 and 2020 in which both acquirer and target are incorporated in the U.S., 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 4-digit code, and equal to 0 otherwise. The variables "Acquirer sales" and "Acquirer employment" are only available for acquirers that can be matched to Compustat. The variable "Number of Patents" shows the number of patents filed before the acquisition date for target firms that have such patents. The variable "Citations (Acquired Patents)" measures the number of citations received by these patents in their initial 5 years.

3.3 The effect of acquisitions on the implementation of ideas

Our model shows that 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 do 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. In our baseline analysis, we use the evolution of citations for patents held by the startup 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. There is empirical support for this proxy: for instance, Argente *et al.* (2025) show that in the consumer goods sector, more highly cited patents lead to a higher likelihood of introducing new products. However, we acknowledge that, like any patent-based measure, it is not perfect. Thus, we will also consider two alternative measures in robustness checks, described below in Section 3.4.

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). To alleviate this concern, we match each acquired startup patent to up to ten control patents. We impose that treatment and control patents have the same patent application year, main technology class and number of inventors, belong to a firm incorporated in the same state and founded in the same year, and are either both above or both below the median of a measure of textual similarity to previous work (developed by Kelly *et al.*, 2021). Online Appendix B.4 discusses the matching variables in greater detail. In particular, Table B.1 shows that there are no significant differences between treated and control patents for other observables as well (including the number of citations received in the patent's first year).

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

$$\text{NumCites}_{ijt} = \beta_1 \text{Acquired}_{ij} + \beta_2 \text{Post}_{jt} + \beta_3 \Big(\text{Acquired}_{ij} \times \text{Post}_{jt} \Big) + \alpha_j + \alpha_t + u_{ijt}, \quad (32)$$

where NumCites_{ijt} is the number of citations received by patent i belonging to matched patent group j in year t,[21] Acquired_{ij} is a dummy taking value 1 if patent i is acquired, and Post_{jt} is a dummy taking value 1 if year t is later than the acquisition of the acquired patent in group j. The baseline specification considers a 10-year window around the acquisition (5 years before and 5 years after this event). Finally, α_j are fixed effects for matched patent groups, and α_t are year fixed effects. In this specification, 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 2 presents the estimation results. Column (1) shows the results of a Poisson estimator without fixed effects. The interaction term is positive, but small and not statistically significantly different from zero. When adding matched group fixed effects (column 2), year fixed effects (column 3), or both sets of fixed effects (column 4), point estimates and significance hardly change. Thus, our results indicate that on average, acquisitions have no effect on startup patent citations. Moreover, note that the estimate for the coefficient of the "Acquired" dummy is always statistically indistinguishable from zero, indicating that treatment and control patents are equally cited before an acquisition. Online Appendix B.5 discusses additional robustness checks, including changes in the number of control patents, different event study windows, and using patent renewals as an implementation proxy.

[21] A matched patent group is defined by the bundle of one treated patent and up to ten control patents. The dependent variable is winsorized at the 99% level.

Table 2: The effect of acquisitions on the implementation of startup ideas

	(1)	(2)	(3)	(4)
Post	1.598***	1.649***	1.295***	0.717***
	(0.143)	(0.129)	(0.159)	(0.112)
Acquired	0.0455	0.0442	0.0707	0.0642
	(0.323)	(0.304)	(0.319)	(0.310)
Acquired × Post	0.0875	0.102	0.0614	0.0803
	(0.291)	(0.283)	(0.288)	(0.287)
Matched Group FE	No	Yes	No	Yes
Year FE	No	No	Yes	Yes
Observations	32,282	28,576	31,610	28,576

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 While the average acquisition has no effect on the implementation of the acquired startup's idea, this could mask underlying heterogeneity. In fact, our model predicts that the effect on implementation depends on the acquisition type, with related acquisitions having a more negative effect than unrelated ones.

To test this prediction, we now estimate a triple-difference specification:

$$\begin{aligned} \text{NumCites}_{ijt} = & \beta_1 \text{Acquired}_{ij} + \beta_2 \text{Post}_{jt} + \beta_3 \text{Related}_{ij} + \beta_4 \Big(\text{Acquired}_{ij} \times \text{Post}_{jt} \Big) \\ & + \beta_5 \Big(\text{Acquired}_{ij} \times \text{Related}_{ij} \Big) + \beta_6 \Big(\text{Post}_{jt} \times \text{Related}_{ij} \Big) \\ & + \beta_7 \Big(\text{Acquired}_{ij} \times \text{Post}_{jt} \times \text{Related}_{ij} \Big) + \alpha_j + \alpha_t + u_{ijt}, \end{aligned} \tag{33}$$

where Related_{ij} is a dummy equal to 1 if the startup owning patent i belongs to the same NAICS 4-digit industry than its acquirer, and to 0 otherwise.

Table 3 summarizes our results, with different columns again corresponding to different combinations of fixed effects. Throughout all specifications, we find that the coefficient on the interaction between the "Acquired" and "Post" dummies is positive and significant at the 10% level. That is, the average unrelated acquisition increases the implementation probability of a startup idea. However, the coefficient on the triple interaction is negative, large, and generally significant at the 5% level, indicating that the average related acquisition lowers the implementation probability of a startup idea. Thus, the zero effect in the overall sample is indeed the average of a negative effect of acquisitions on citations for related acquisitions, and a positive effect for unrelated acquisitions.

Table 3: Heterogeneous effects: related and unrelated acquisitions

	(1)	(2)	(3)	(4)
Post	1.457***	1.491***	1.153***	0.607***
	(0.174)	(0.159)	(0.183)	(0.132)
Acquired	-0.432	-0.361	-0.399	-0.357
	(0.359)	(0.351)	(0.350)	(0.349)
Acquired × Post	0.557*	0.522*	0.525*	0.518*
	(0.303)	(0.306)	(0.294)	(0.297)
Related	-0.754**		-0.781**	
	(0.313)		(0.376)	
Acquired × Related	1.198**	1.005*	1.180*	1.026*
	(0.606)	(0.577)	(0.610)	(0.597)
Post × Related	0.461*	0.504**	0.489	0.361
	(0.269)	(0.239)	(0.301)	(0.220)
Acquired × Post × Related	-1.161**	-1.046**	-1.150**	-1.069*
	(0.537)	(0.528)	(0.540)	(0.548)
Matched Group FE	No	Yes	No	Yes
Year FE	No	No	Yes	Yes
Observations	32,282	28,576	31,610	28,576

Notes: We use a Poisson estimator. The dependent variable is the number of citations received at the patent-year level. The dummy variables "Acquired" and "Post" are defined as in Table 2. "Related" is a dummy equal to 1 if the acquired startup and the acquirer belong to the same NAICS 4-digit industry. Standard errors are clustered at the target firm level. * significant at 10%; ** significant at 5%; *** significant at 1%.

3.4 An alternative implementation measure: drug development

As discussed, patent citations might not be a perfect proxy for idea development. Thus, we now complement our evidence by studying a sample of drug development projects by pharmaceutical companies, for whom we have more precise implementation measures.

Data Our drug development dataset was introduced by Cunningham *et al.* (2021) in their "killer acquisitions" paper. The primary data source is Pharmaprojects, a database compiled by the private company Citeline, which tracks the history of drug development projects globally. Pharmaprojects identifies a number of "development events" for each drug (e.g. new launches, approvals by health agencies, etc.), which can be used as a proxy for implementation. We match this database to SDC Platinum in order to identify pharmaceutical companies that were acquired. Our final sample contains 15,023 drugs which had at least one development event between 1989 and 2010, and 255 acquired firms for whom we know the incorporation date. Online Appendix B.6 describes the data construction, which closely follows Cunningham *et al.* (2021).

Empirical Methodology As before, we investigate the effect of acquisitions on startup ideas. Due to a lower sample size, we now define startups as targets that were below the median target age at acquisition. We then estimate a difference-in-difference specification:

$$\text{Development}_{it} = \beta_1 \text{Acquired}_i + \beta_2 \Big(\text{Acquired}_i \times \text{Post}_{it} \Big) + \alpha_{FE} + u_{it}, \tag{34}$$

where Development_{it} is a dummy variable equal to 1 when drug i has a development event in year t. Acquired_i is a dummy variable for drug i belonging to a startup that is acquired during the sample period, and Post_{it} indicates whether the drug-year observation (i, t) is posterior to the acquisition. α_{FE} stands for different fixed effects, described below. Standard errors are clustered at the drug project level.

Table 4: The effect of acquisitions on startup drug projects

	(1)	(2)	(3)	(4)
Acquired	0.00956	0.0155		
	(0.0146)	(0.0148)		
Acquired × Post	-0.0154	-0.0299	-0.0253	-0.0598
	(0.0201)	(0.0202)	(0.0274)	(0.0375)
Age FE	No	Yes	Yes	Yes
Vintage FE	No	Yes	Yes	Yes
Project FE	No	No	Yes	Yes
Age × TC × MOA	No	No	No	Yes
Observations	102,563	102,563	102,474	74,988

Notes: We use a OLS estimator. Development_{it} is a dummy variable taking value 1 when a drug i has a development event in year t. Acquired_i is a dummy equal to 1 if drug i belongs to an acquired startup, and Post_{it} indicates whether the drug-year observation (i, t) is posterior to the acquisition. Standard errors are clustered at the drug project level. * significant at 10%; ** significant at 5%; *** significant at 1%.

Findings Table 4 shows our results. Column (1) does not include any fixed effects. Column (2) includes fixed effects for the age of the drug project and for its vintage (the year in which it started). Column (3) adds drug-level fixed effects, and Column (4) interacts age fixed effects with fixed effects for a drug's therapeutic class (TC) and mechanism of action (MOA), two key drug characteristics. In all specifications, we find that the coefficient on the interaction Acquired $\times$ Post is indistinguishable from zero: on average, an acquisition does not affect the development probability of a startup's drug project.

Note that this result does not contradict the "killer acquisition" finding of Cunningham *et al.* (2021). Indeed, their paper estimates a triple-difference version of equation (34) and finds that acquisitions lower the implementation probability only for drug projects that overlap with those of the acquirer. In Online Appendix B.6, we show that we can replicate this result. This very much echoes our own previous finding in Table 3, namely that only related acquisitions lower the implementation probability.

Summing up, our empirical results suggest that a ban on startup acquisitions would not lead to more implemented startup ideas, but a ban on related acquisitions might. However, this preliminary conclusion could well be misleading. First, our difference-in-difference regressions miss general equilibrium effects that affect all firms equally. Second, as our model 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.

4 Calibration

4.1 Externally calibrated parameters

We consider a BGP without acquisition taxes ($\tau_R = \tau_U = 0$), where a period of unit length corresponds to one year. We set the discount rate to $\rho = 0.03$, which together with a 2% target for growth implies a real interest rate of 5%. There are two values for product spending shares, ω_L and ω_H, with $\omega_L < \omega_H$. 10% of products belong to the H class, and we set $\omega_H/\omega_L = 16.125$ following Hottman *et al.* (2016), who show that for 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 calibrated: 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, ζ_I^{int} and ζ_I^{ext}; the startup creation cost, ζ_S; and the implementation cost shifters for incumbents and startups, κ_I and κ_S.

We calibrate these parameters through indirect inference, choosing parameter values to minimize the distance between ten model-generated moments and their empirical counterparts. As the model is non-linear, each moment is a priori affected by all parameters, making identification challenging. Our discussion in the main text stresses economic intuitions for identification, while Online Appendix C.3 shows the results of a formal identification test, indicating that parameters are well identified by the chosen moments.

4.2.1 Startup acquisition moments

We target several moments from our dataset and empirical analysis. First, our regressions in Section 3.3 showed that the average acquisition has no significant effect 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. This moment identifies the relative implementation cost of incumbents, κ_I/κ_S. As shown in Table 5, we find that incumbent implementation costs are 29% lower than those of startups. This compensates for the replacement effect which, all else equal, would have implied a negative coefficient in the model regression.

Second, we target a share of unrelated acquisitions of 60.9%, corresponding to the share of startup acquisitions in which target and acquirer belong to different NAICS 4-digit industries (see Table 1). This moment pins down γ, the probability that a startup idea is matched to an unrelated incumbent.

Third, we target an average acquisition premium of 50.7% (see Table 1). In the model, the acquisition premium is the difference between the acquisition price and the startup's outside option. This moment pins down α, the incumbent's bargaining weight. We find $\alpha = 0.668$, implying that incumbents obtain a larger surplus share than startups.

Fourth, we target the average frequency of acquisitions for incumbents. This moment identifies χ, which shifts search costs and therefore governs the probability with which incumbents notice and acquire startups.[22] In the data, we use our matched SDC-Compustat dataset, and consider Compustat firms with positive R&D spending as the equivalent of

[22]To solve 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 distribution of spending shares and technology gaps across firms. Online Appendix C.2 contains further details.

incumbent firms in our model.[23] Obviously, these firms are not representative for the population of U.S. firms. However, calibrating our endogenous growth model to the average firm would be misleading: in our model, all incumbents spend on R&D, while in the data, this is only true for a small minority of firms. Thus, we calibrate the model to the population of innovating firms (a common practice in the literature, see e.g. Akcigit and Kerr, 2018), and use Compustat firms with positive R&D spending as a proxy for these innovating firms.

Computing the average yearly number of startup acquisitions made by Compustat firms with positive R&D spending, we find an acquisition frequency of 3.17%. However, this number is a lower bound. As discussed in Section 3, SDC Platinum does not always provide a founding date for acquisition targets: for 39.4% of acquisitions made by Compustat firms with positive R&D spending, the founding date is missing. If the age distribution among these targets was the same as the age distribution among the targets with an observed founding year, the true startup acquisition frequency would be 5.23% ($\frac{0.0317}{1-0.394}$). To be conservative, we target the midpoint between this estimate and our lower bound, a startup acquisition frequency of 4.20%. However, we show in Section 5.4 that our policy results are unchanged for alternative values of this target, ranging between 3 and 6%.

Fifth, our model implies that larger firms are more likely to acquire startups. As this is also the case in the data, we aim to quantitatively match this relationship. To do so, we run a simple regression of the number of startup acquisitions on firm size (as measured by employment) for our sample of Compustat firms with positive R&D spending. We estimate

$$\text{SAQ}^s_{f,t} = \alpha^s_t + \beta_S \ln\left(\text{Emp}^s_{f,t}\right) + \varepsilon^s_{f,t}, \tag{35}$$

where $\text{SAQ}^s_{f,t}$ is the number of startups acquired by firm f of industry s in year t, $\text{Emp}^s_{f,t}$ is the number of employees of firm f in year t, and α^s_t is a set of industry-year fixed effects, with industries defined at the 4-digit NAICS level. We normalize the dependent variable to have mean 1, and find $\widehat{\beta}_S = 0.558$ (with a 95% confidence interval between 0.517 and 0.599).[24] That is, there is a strong positive link between firm size and startup acquisitions. We target the same regression result in our model. This moment identifies φ, the curvature of search costs: the lower the curvature, the stronger the link between acquisitions and size.

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

[23] We consider the period 2000-2020 (as in our empirical analysis), and limit the Compustat sample to firms that spend on R&D in every year in which they are observed.

[24] Normalizing is crucial for robustness checks in which we vary the frequency of startup acquisitions. Without normalization, this change in the scale of the left-hand side variable would mechanically change the coefficient β as well. Table B.3 in the Online Appendix shows the full results for this regression.

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 Online Appendix C.4, this target matters little for the optimal policy.

4.2.2 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), which identifies the innovation step size λ. Moreover, 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 78.9% to incumbents. We target the same shares in the model. This identifies the level of incumbent innovation costs, ξ_I^{int} and ξ_I^{ext}, which determines incumbents' growth contribution. 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, $\xi_I^{\text{ext}}/\xi_I^{\text{int}}$. 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.[25] This identifies the startup creation cost ξ_S, which through the free entry condition shifts the startup rate and hence the exit rate of incumbents.

4.2.3 Model fit and untargeted moments

Table 5 lists the calibrated parameter values and shows that our model fits the targeted moments very closely.[26]

[25] This is an approximation, as large firms in Garcia-Macia *et al.* (2019) do not correspond to Compustat firms with positive R&D spending. However, other studies imply similar values. 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.

[26] Online Appendix C.1 describes how we compute model moments. The vector of internal parameters $\boldsymbol{\theta}$ is chosen to minimize $\sum_{m=1}^{10} \widetilde{w}_m \frac{|\text{Moment}_m(\text{Model},\boldsymbol{\theta})-\text{Moment}_m(\text{Data})|}{0.5|\text{Moment}_m(\text{Model},\boldsymbol{\theta})|+0.5|\text{Moment}_m(\text{Data})|}$, where $\widetilde{w}_m$ captures the weight on moment m. The only role of the weights is to ensure that the growth rate target is matched exactly: the weight on this target is set to 10, and the weight on all other targets to 1.

Table 5: Parameter values and model fit in the baseline calibration

Parameter	Value	Interpretation	Source / Moment	Data	Model
ρ	0.03	Discount rate	Standard value		
ϕ_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	1.042	Innovation step size	Growth rate	2.00%	2.00%
χ	0.348	Search cost scale	Frequency of startup acq.	4.20%	4.19%
κ_I/κ_S	0.713	Incumbent relative impl. cost	Reg., implem. on acq.	0.000	-0.000
α	0.668	Incumbent bargaining weight	Acquisition premium	50.7%	51.0%
φ	2.984	Search cost curvature	Reg., acq. on firm sales	0.558	0.558
γ	0.659	Share of unrelated startups	Share of unrelated acq.	60.9%	60.8%
ξ_I^{int}	0.94	Incumbent int. innovation cost	Entry contrib. to growth	21.1%	21.1%
ξ_I^{ext}/ξ_I^{int}	6.356	Rel. external innovation cost	Incumbent CD share	12.6%	12.6%
ξ_S	0.063	Startup cost	Exit rate	6.0%	6.0%
κ_S	2.414	Startup implementation cost	Average impl probability	10.0%	10.0%

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 its data and model values. The abbreviation "Reg." stands for regression coefficient.

Our calibrated model is also in line with several untargeted moments. Figure 4 compares the employment distribution of incumbent firms in the model to the employment distribution of Compustat firms with positive R&D spending, by plotting the Lorenz curve of employment for both model and data.[27] The firm employment distribution in our model is strongly skewed, though slightly less than in the data.

Figure 4: The Lorenz curve of firm employment in the model and in the data

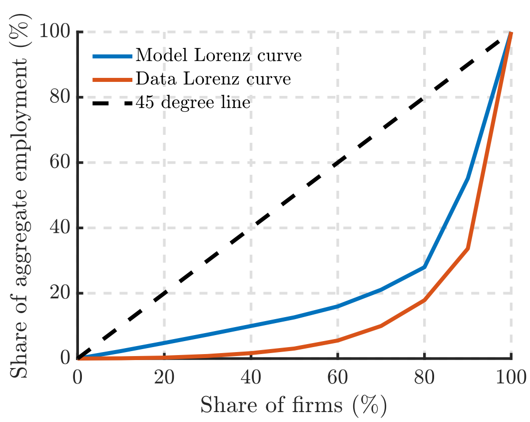

[27]In the data, we first compute for each NAICS 4-digit industry the share of total employment for each decile of the industry employment distribution, using data from 2010. We then take a weighted average of these values, where the weights are total industry employment. In the model, we compute the employment distribution of firms using the same simulated data that we use to compute some of our targeted moments.

As shown previously in Figure 2, our calibration is also qualitatively in line with our regression results in Table 3, indicating that related acquisitions lower and unrelated acquisitions increase the implementation probability of a startup idea.

5 Regulating startup acquisitions

5.1 Some primers for policy analysis

We are now ready to analyze the economic effect of acquisition policies, implemented through changes in the acquisition taxes τ_R and τ_U. To do so, we assume that the economy is initially on the no-policy BGP, when the government makes a surprise announcement of a new policy (i.e., a new level of acquisition taxes) at $t = 0$. We then solve for the transition path to the new BGP implied by this policy. Computing this transition path is not straightforward, as it involves iterating over the path of aggregate interest and growth rates. Online Appendix A.6 describes our algorithm for solving this problem.

To assess the welfare implications of a given policy, we compare the path of consumption over the transition to a counterfactual path of consumption that would have prevailed if the economy had remained in the no-policy BGP. The consumption-equivalent welfare change γ^{CEW} triggered by the policy is then pinned down by the condition

$$\int_0^{+\infty} \exp(-\rho t) \ln\left(\gamma^{\text{CEW}} C_t^{\text{No Policy}}\right) dt = \int_0^{+\infty} \exp(-\rho t) \ln\left(C_t^{\text{Policy}}\right) dt, \tag{36}$$

where $C_t^{\text{No Policy}}$ stands for consumption at time t in the no-policy BGP, and C_t^{Policy} for consumption at time t with the new policy. The consumption-equivalent welfare change measures the permanent shift in consumption in the no-policy BGP that would make the consumer indifferent between this BGP and the transition to the new policy BGP.

Throughout our analysis, the most important driver of welfare is the change in the long-run growth rate induced by different policies. To understand this change, we now introduce a useful decomposition result. Online Appendix A.7 shows that the change in the growth rate between the no-policy BGP and a BGP with a policy can be expressed as:

$$\begin{aligned}\frac{g^{\text{Policy}}}{g^{\text{No Policy}}} = {} & \text{Share}_{\text{inc}}^{\text{No Policy}} \cdot \frac{\text{Incumbent own innovation}^{\text{Policy}}}{\text{Incumbent own innovation}^{\text{No Policy}}} \\ & + (1 - \text{Share}_{\text{inc}}^{\text{No Policy}}) \cdot \left(\frac{\text{Startup rate}^{\text{Policy}}}{\text{Startup rate}^{\text{No Policy}}} \cdot \frac{\text{Impl. rate of startup ideas}^{\text{Policy}}}{\text{Impl. rate of startup ideas}^{\text{No Policy}}}\right),\end{aligned} \tag{37}$$

with

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

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

$$\text{Impl. rate of startup ideas} = \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m(\omega,a)\omega i(\omega,a).$$

Thus, the change in long-run growth is a weighted average of three margins: (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). The weight on incumbents' own innovation is the percentage of growth in the no-policy BGP explained by this margin, around 76% in our baseline calibration.

5.2 Should killer acquisitions be banned?

The public debate on startup acquisitions has largely been shaped by concerns about killer acquisitions. Therefore, we first evaluate a simple policy which bans all killer acquisitions (i.e., all acquisitions that lower the implementation probability of startup ideas). In our baseline calibration, this is equivalent to banning all related acquisitions by imposing an infinitely high tax τ_R, while leaving unrelated acquisitions untouched.[28]

If such a policy would only reshuffle the implementation probabilities of ideas, it would indeed mechanically increase growth and welfare. To get a sense of magnitudes, we compute a counterfactual path of consumption obtained when setting the frequency of killer acquisitions to zero (i.e., $s_R(\omega, a) = 0$) upon the announcement of the policy, but leaving all other innovation and acquisition decisions (the startup rate, incumbents' own innovation, search efforts for unrelated acquisitions, etc.) unchanged.[29] We find that in this case, the ban raises the long-run growth rate by 4 basis points, and generates a significant consumption-equivalent welfare gain of 1.29%.

However, our model suggests that this partial equilibrium assessment could be misleading: a ban on killer acquisitions clearly affects the incentives for creating a startup or investing in R&D. How do these general equilibrium feedback effects change the picture?

To answer this question, we now turn to the full transition outcomes, allowing all firm decisions to adjust endogenously to the announcement of the policy. Figure 5 summarizes

[28]Figure 2 shows that all related acquisitions lower the implementation probability of a startup idea. Interestingly, this policy is in line with the Platform Competition and Opportunity Act discussed earlier, which would prohibit all acquisitions of direct, nascent or potential competitors for Technology Platforms.

[29]We still compute a full transition path in this case, using the counterfactual policy functions.

our results, by plotting the response of output growth (Panel A) and the consumption share of output (Panel B). These two variables fully pin down the policy's effect on aggregate consumption. As the figure shows, results are very different from the partial equilibrium analysis. Output growth now drops on impact, and long-run growth settles at 1.985% per year: a 1.5 basis point drop instead of a 4 basis points increase. As the policy increases the consumption share of output by 0.4 percentage points, it still leads to a tiny overall welfare gain of 0.04%.[30] However, as this gain is not meaningfully different from zero, we can conclude that general equilibrium feedback effects eliminate virtually all of the direct gains from a killer acquisition ban.

Figure 5: The path of output growth and the consumption share after a ban on related acquisitions

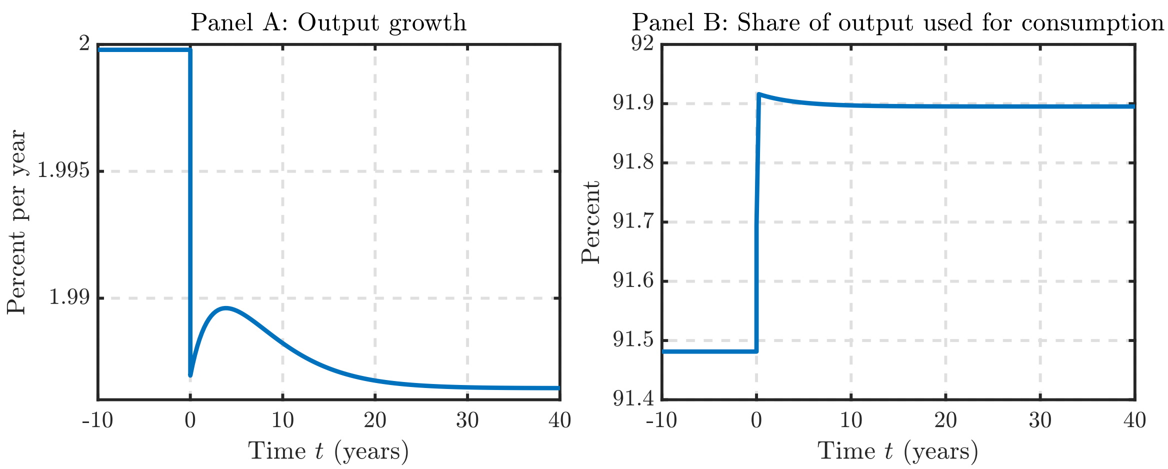

Notes: This figure illustrates the transition from a no-policy BGP to a BGP with a ban on related acquisitions. We assume that the economy is initially in the baseline (no-policy) BGP, and the government announces at time $t = 0$ an infinitely high tax τ_R on related acquisitions. Panel A shows the path of the (annualized) rate of output growth over the transition to the new BGP, while Panel B shows the path of the share of output used for consumption. The algorithm used to solve for the transition is described in Online Appendix A.6.

What drives this attenuation effect? Table 6 provides further details by listing the equilibrium values of several key variables in the initial BGP (column (A)) and the new BGP implied by the ban (column (B)). In particular, we can now use our decomposition from equation (37) to analyze changes in long-run growth. The small drop in overall growth masks large changes in the underlying margins. The killer acquisition ban lowers the startup rate substantially, by 16.5%.[31] This is only partially offset by a 1.2% increase in the own

[30] For both variables, the transition to the new BGP is almost immediate. This is due to the block nature of the model, which implies that innovation rates, implementation probabilities and search efforts essentially jump to their new BGP levels. Transition dynamics are only due to composition effects through the slow-moving technology gap distribution, and these turn out to be small.

[31] 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%.

innovation of incumbents and a 11.9% increase in the implementation rate of startup ideas.

Table 6: The effects of acquisition taxes on growth and welfare

	(A) No Policy	(B) Related ban	(C) Complete Ban	(D) Optimal Policy
Policy				
Related acq. tax (τ_R)	0%	$+\infty$	$+\infty$	215%
Unrelated acq. tax (τ_U)	0%	0%	$+\infty$	30%
BGP outcomes				
Growth rate	2.00%	1.99%	1.99%	2.01%
Startup acquisition frequency	4.19%	2.17%	0.00%	1.59%
Related acq. frequency	1.64%	0.00%	0.00%	0.49%
Unrelated acq. frequency	2.55%	2.17%	0.00%	1.10%
Incumbent own innovation	1	1.012	1.013	1.026
Startup rate	1	0.835	0.842	0.864
Startup idea implementation	1	1.119	1.105	1.077
Consumption share	91.5%	91.9%	91.9%	91.7%
Misallocation loss	0.49%	0.42%	0.43%	0.53%
Transition dynamics				
CE welfare change		0.04%	0.03%	0.48%

Notes: In this table, each column lists key BGP outcomes for a given policy on startup acquisitions. The three growth margins (incumbent own innovation, the startup rate and startup idea implementation) are expressed relative to their value in the no-policy BGP. Consumption-equivalent ("CE") welfare changes are computed as described in equation (36).

To understand these changes, recall that startups agree to acquisitions because these generate a surplus over their outside option. In particular, in killer acquisitions, incumbents are forced to share some of their rents with the startup. Banning these operations eliminates that surplus and lowers the incentives for startup creation. The ensuing lower startup rate has a direct negative effect on growth, but it also has wide-ranging and ambiguous general-equilibrium effects on incumbents.

On the one hand, there is a direct negative effect: the ability to acquire related startups created a surplus for incumbents, which is eliminated by the ban. All else equal, this lowers incumbent value and innovation incentives. On the other hand, the lower startup rate implies that incumbents are less likely to be displaced or threatened by startups, increasing their value and innovation incentives. In our calibration, this positive effect slightly dominates, explaining an increase in both incumbents' own innovation and in the effort of startups to implement their ideas (since startups aim to become incumbents). Together with the mechanical increase in the implementation probability from the ban of killer acquisitions, this explains the increase in the startup idea implementation margin.

While output growth is the most important driver of welfare, the share of output devoted to consumption matters as well. The killer acquisition ban increases this share by 0.4 percentage points, maintaining a tiny overall welfare gain.[32] This is due to a decrease in the share of output used for startup creation (falling by 0.30 percentage points) and for startup search (falling by 0.20 percentage points), dampened by a slight increase in the share of output used for incumbent innovation (0.09 percentage points).[33]

Our results so far suggest that banning killer acquisitions does not significantly increase welfare. However, are there other policies that could achieve this goal?

5.3 Optimal startup acquisition taxes

To answer this question, we compare the economy's response to a policy shock across a large grid of potential values for acquisition taxes τ_R and τ_U. We find that the optimal policy (i.e., the policy that maximizes consumption-equivalent welfare) imposes a high 215% tax on related acquisitions, and a lower 30% tax on unrelated acquisitions. As shown in Column (D) of Table 6, this policy reduces the frequency of related acquisitions by roughly 70%, and the frequency of unrelated acquisitions by more than 50%. It triggers a 1 basis point increase in long-run growth, which together with a 0.2 percentage point increase in the consumption share yields a 0.48% increase in consumption-equivalent welfare.

To understand why this policy is optimal, Figure 6 plots several key outcomes as a function of the related acquisition tax τ_R, keeping the unrelated acquisition tax at its optimal level of 30% throughout. Indeed, the general equilibrium effects of unrelated acquisitions are minor, and our main results are driven by related acquisitions.

Panel A reveals an inverted U-shape: growth and welfare first increase and then decrease in the tax rate τ_R. Panel B indicates that this pattern is driven by incumbent own innovation, which also has an inverted U-shape in the tax rate. In our calibration, incumbent own innovation is the main driver of growth, accounting for around three quarters of the aggregate growth rate. Hence, the optimal tax essentially maximizes the positive response of incumbent own innovation, lifting it by 2.6% (rather than 1.2% with the killer acquisition ban). This has the positive side effect of also achieving a smaller fall in the startup rate (13.6%, rather than 16.5% with the ban).

To understand this U-shape, recall that the tax has an ambiguous effect on incumbent value: there is a direct negative effect (through lower acquisition surpluses) and an indirect

[32] As noted in the introduction, policy-induced changes in static misallocation of labor across firms are small and have a negligible effect on welfare, so we disregard them in this discussion.

[33] Banning all startup acquisitions would have very similar effects to a killer acquisition ban, as shown in column (C) of Table 6. As unrelated acquisitions generate only a small surplus, banning them does not trigger large general equilibrium spillovers.

positive effect (through a lower startup rate). Crucially, these effects are not uniform across the technology gap distribution. Gains from a lower startup rate do not depend on the technology gap. However, losses from lower acquisition surpluses affect mostly low-technology-gap products, for which related acquisitions are no longer profitable. For high-technology-gap products, there is still a surplus, and the acquirer's share of this surplus actually increases as a function of the tax. This is shown in the left panel of Figure 7, and immediately apparent from the expression of the acquirer's surplus share in equation (23). Indeed, as the tax-induced loss in joint surplus is increasing in the acquisition price, the tax pushes the price down, benefiting the acquirer.[34]

Figure 6: The effect of related acquisition taxes on growth and welfare.

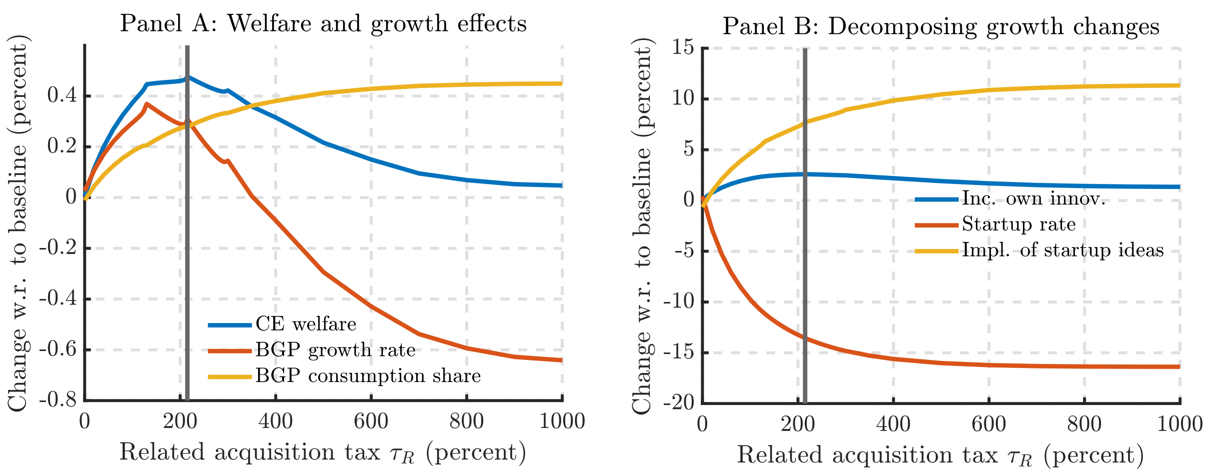

Notes: This figure illustrates the impact of different levels of the related acquisition tax τ_R. We assume that the government announces a new level of this tax (and a 30% tax on unrelated acquisitions) at time $t = 0$, and then solve for the transition to the new BGP. Panel A shows the growth rate and consumption share in the new BGP for each value of τ_R, as well as the consumption-equivalent welfare change from the transition. The gray vertical line indicates the optimal related acquisition tax. Panel B decomposes the change in the BGP growth rate into the three margins discussed in Section 5.1.

The right panel of Figure 7 illustrates the impact on product values, by plotting the product-level value function in the no-policy and optimal policy BGPs. As high-technology-gap products are shielded from the worst effects of the tax, their value increases and the value function becomes steeper in the optimal policy BGP. Thus, incumbents have more incentives to increase their technology gaps through innovation, in order to be able to protect themselves again through related acquisitions. In contrast, a ban on related acquisitions would eliminate this incentive, and the slope of the value function would remain unchanged.

[34]This is a well-known result in the literature using Nash bargaining (see e.g. Han *et al.* (2025) for a recent example in the context of the housing market).

The optimal policy also imposes a small tax on unrelated acquisitions. Unrelated acquisitions boost the implementation probability of startup ideas, and therefore increase creative destruction. The optimal policy curbs this threat somewhat, which again shores up incumbent values and stimulates incumbent own innovation.

Figure 7: The impact of acquisition taxes on surplus shares and product values

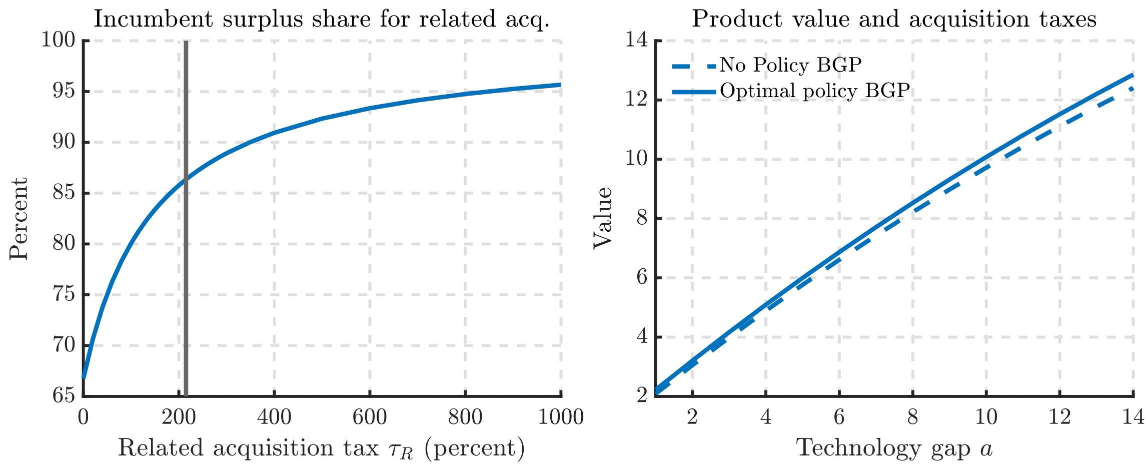

Notes: The left panel of this figure plots the incumbent's surplus share from a related acquisition, $\widetilde{\alpha}_R$, as a function of the related acquisition tax τ_R. The right panel shows the value function v for high spending share products in the initial no-policy BGP (dashed line) and the BGP with optimal acquisition taxes (solid line).

5.4 Robustness checks

The previous section has shown that general equilibrium effects are crucial for designing policy on startup acquisitions. However, our quantitative results are not hard-wired features of the model, but depend on the data we use to calibrate it. In this section, we explore the role of these calibration targets and consider several model extensions.

5.4.1 Calibration targets

First, we consider the role of some important calibration targets for our results, starting with the acquisition premium in Figure 8. In our baseline calibration, the acquisition premium is 50.7%, as indicated by the vertical lines in the figure. Each point on the x-axis shows an alternative calibration target for the premium, leaving all other targets and all external parameters unchanged. The different panels show the optimal acquisition taxes, the reduction in the frequency of acquisitions, and the consumption-equivalent welfare gains from the optimal policy in these re-calibrated models.

As the figure shows, higher acquisition premia make bans more desirable: while the optimal policy rates hardly change, the consumption-equivalent welfare gain increases linearly in the acquisition premium. Indeed, this target disciplines the costs and benefits of acquisitions. High premia imply that acquisitions are lucrative for startups, but costly for incumbents. Thus, with high premia, taxing 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 startups). As incumbents contribute more to aggregate growth than startups, the latter effect dominates.

Figure 8: Robustness to different targets for the acquisition premium.

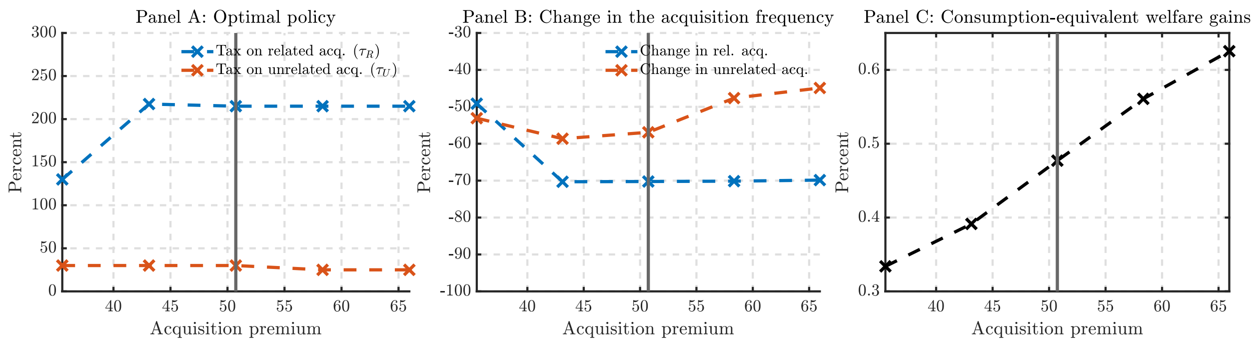

Notes: Each point on the x-axis corresponds to a re-calibrated model with a different target for the acquisition premium. All other calibration targets and external parameters are at their baseline values.

Figure 9 examines the role of our estimates for the effect of acquisitions on the citations to pre-existing startup patents, our proxy for idea implementation.

Figure 9: Robustness to different targets for the effect of acquisitions on the implementation of startup ideas.

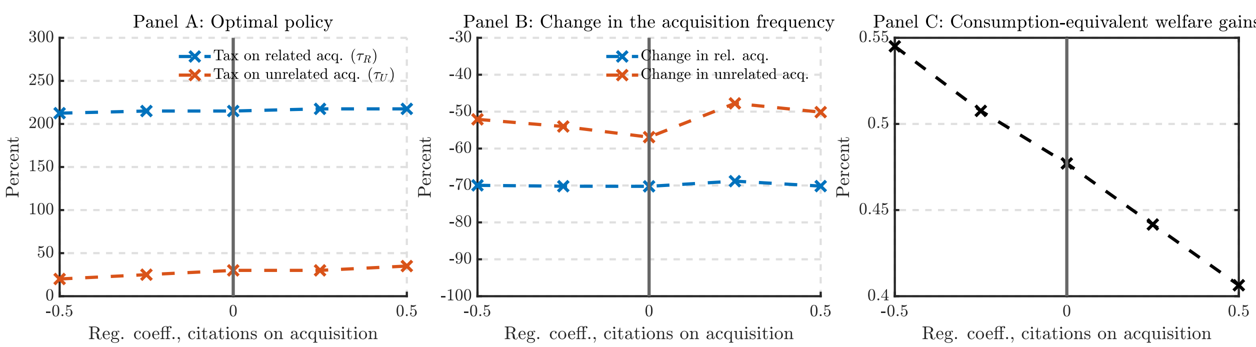

Notes: Each point on the x-axis corresponds to a re-calibrated model with a different target for the effect of acquisitions on patent citations. All other calibration targets and external parameters are at their baseline values.

The baseline target for this number was zero. As the figure shows, optimal tax rates are invariant in the targeted regression coefficient, but welfare gains are decreasing: taxing acquisitions is more desirable when killer acquisitions are more prevalent.[35] However,

[35] In order to target non-zero coefficients, we need to assume an elasticity of the implementation probability

comparing the scale of the x-axis with the magnitude of our empirical estimates in Table 2 shows that these differences are small, and our regressions would have to yield considerably higher numbers to move the needle.

Figure 10: Robustness to different targets for the startup acquisition frequency.

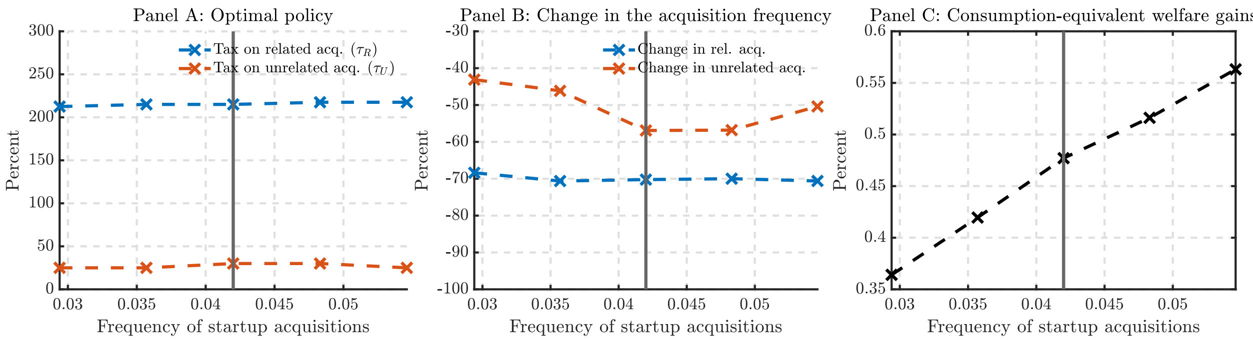

Notes: Each point on the x-axis corresponds to a re-calibrated model with a different target for frequency of startup acquisitions. All other calibration targets and external parameters are at their baseline values.

Finally, Figure 10 shows how our results change for different targets for the frequency of startup acquisitions, a moment which is only observed imperfectly in the data. Optimal acquisition taxes are virtually unchanged for different levels of this target, but welfare gains are increasing slightly in the frequency of startup acquisitions: startup acquisitions matter more when they are more prevalent.

In Online Appendix C.4, we expand on the analysis in this section by showing the results of robustness checks with respect to all other calibration targets.

5.4.2 Alternative free entry assumptions

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 (e.g. Klette and Kortum, 2004; Peters, 2020), it obviously influences the strength of the response of the startup rate to any policy.

To assess the robustness of our results, we now relax this assumption, and assume that that startup creation costs are given by $\xi_S (x_S)^\nu Y_t$. In our baseline model, $\nu = 0$. Instead, allowing $\nu > 0$ generates congestion effects, with the cost of startup creation increasing in

of ideas (the outcome in our model) to patent citations (the outcome in our empirical analysis). We rely on Kogan *et al.* (2017), who estimate that the elasticity of a patent's market value to its number of forward citations is 0.174. In our model, the expected value of an idea is equal to the product of the implementation probability and the value of the implemented idea, and the latter is ex ante identical for all startup ideas. Thus, we use this elasticity for our robustness checks.

the mass of startups. We consider values of $\nu = 0.14$ (in line with the estimation results of Klenow and Li, 2025) and $\nu = 1$. We then re-calibrate the model to match the baseline targets, leaving all other external parameters unchanged. Table 7 shows that congestion has only minor effects: optimal tax rates and consumption-equivalent welfare gains are virtually unchanged. However, the policy now triggers a smaller fall in the startup rate.

Table 7: Robustness checks: external parameters

Parameter values	τ_R	τ_U	Change in CEW	Startup acq.	Startup rate
Baseline	215.0%	30%	0.48%	-62.13%	-13.55%
$\nu = 0.14$	217.5%	30%	0.46%	-61.05%	-11.72%
$\nu = 1$	217.5%	25%	0.55%	-53.35%	-6.84%
$\frac{\omega_H}{\omega_L} = 30$	142.5%	30%	0.64%	-58.97%	-16.36%

Notes: The first row of this table recalls the results for our baseline calibration. In the second and third row, we change the value of the congestion parameter ν, and in the fourth row, we change the value for the relative spending share of high-spending-share products, $\frac{\omega_H}{\omega_L}$. In each case, we re-calibrate the model, leaving all targeted moments and all other external parameters at their baseline values. The first two columns list optimal acquisition taxes, and the last three columns show the changes in key variables due to this policy. The abbreviation “CEW” stands for consumption-equivalent welfare.

5.4.3 Matching the firm size distribution

As shown in Figure 4, our model does not perfectly match the employment distribution of innovating incumbents (a non-targeted moment in our calibration). To address this, we consider a robustness check in which we change the relative spending share of high-spending-share products (i.e., the ratio $\frac{\omega_H}{\omega_L}$) to match the eight decile of the Lorenz curve of employment (i.e., the employment share of the bottom 80% of firms). This implies a relative spending share of 30, as opposed to 16.125 in the baseline.[36]

As shown in Table 7, our policy conclusions in this alternative calibration are close to the baseline. Again, the optimal policy imposes a high tax on related acquisitions (of 142.5%), and a lower tax on unrelated acquisitions (of 30%). The consumption-equivalent welfare gains in this case are somewhat larger than in the baseline, but of a comparable magnitude.

5.4.4 Model extensions

Online Appendix D discusses two extended versions of our baseline model. These allow for quality differences between startup and incumbent ideas, or enable incumbents to

[36] As in every other robustness exercise, we re-calibrate all internal parameters to match our baseline targets. In this calibration, the employment share of the bottom 80% of firms is 18%, as in the data.

acquire ideas from other incumbents. In both cases, our main results remain unchanged.

All in all, these robustness checks show that our quantitative results are robust to reasonable variations in calibration targets and model assumptions. However, they also indicate that welfare gains depend on the characteristics of the economy and could therefore vary across different countries, time periods, or industries.

6 Conclusion

Startup acquisitions have ambiguous effects on innovation and growth. In this paper, we have used a general equilibrium model that takes into account several positive and negative effects of acquisitions, and disciplined it by calibrating it to micro-level data. Despite the intuitive appeal of a ban on killer acquisitions, we find general equilibrium effects eliminate the gains from such a policy, mostly due to a large drop in startup creation. Policy makers can obtain much greater welfare gains by imposing substantial taxes on startup acquisitions. These reduce acquisitions to the most valuable products, maximize incumbent innovation incentives, and increase consumption-equivalent welfare by 0.48%.

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, welfare gains from acquisition taxes are higher when acquisition premia are high and killer acquisitions are frequent. Exploring this heterogeneity more systematically is a promising path for future research. Likewise, future research could speak to further ways in which acquisitions might affect innovation, such as the reallocation of researchers across research teams and firms.

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. (2025). *Patents to Products: Product Innovation and Firm Dynamics*. Working Paper 34592, NBER.

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*. Working Paper 22750, NBER.

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*. Discussion Paper 13115, CEPR.

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. (2025). Mergers and Acquisitions and the Aggregate Markup. *Macroeconomic Dynamics*, **29**, e115.

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. (2026). The Macroeconomic Consequences of Early Childhood Development Policies. *Journal of Political Economy*, forthcoming.

David, J. M. (2020). The Aggregate Implications of Mergers and Acquisitions. *The Review of Economic Studies*, **88** (4), 1796–1830.

De Haas, R., Sterk, V. and van Horen, N. (2022). *Start-up Types and Macroeconomic Performance in Europe*. Staff Working Paper No. 986, Bank of England.

Decker, R., Haltiwanger, J., Jarmin, R. and Miranda, J. (2014). The Role of Entrepreneurship in US Job Creation and Economic Dynamism. *Journal of Economic Perspectives*, **28** (3), 3–24.

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. *Mimeo*.

EWENS, M. and MARX, M. (2024). Firm Age and Invention: An Open-Access Dataset. *Mimeo*.

—, PETERS, R. and WANG, S. (2025). Measuring intangible capital with market prices. *Management Science*.

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.

GUTHMANN, R. R. and RAHMAN, D. M. (2025). Competition Policy in Light of Innovation. *Mimeo*.

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.

HAN, L., NGAI, L. R. and SHEEDY, K. D. (2025). To Own or to Rent? The Effects of Transaction Taxes on Housing Markets. *The Review of Economic Studies*, p. rdaf092.

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.

IGNASZAK, M. and SEDLÁČEK, P. (2023). *Customer Acquisition, Business Dynamism and Aggregate Growth*. Discussion Paper 16205, CEPR.

JONES, C. I. (2016). The Facts of Economic Growth. In J. B. Taylor and H. Uhlig (eds.), *Handbook of Macroeconomics*, vol. 2, Elsevier, pp. 3–69.

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

KAMEPALLI, S. K., RAJAN, R. and ZINGALES, L. (2020). *Kill Zone*. Working Paper 27146, NBER.

KELLY, B., PAPANIKOLAOU, D., SERU, A. and TADDY, M. (2021). Measuring Technological Innovation over the Long Run. *American Economic Review: Insights*, **3** (3), 303–20.

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

KLENOW, P. J. and LI, H. (2025). Entry Costs Rise with Growth. *Journal of Political Economy Macroeconomics*, **3** (1), 43–74.

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*. Working Paper 7345, NBER.

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. (2024). Killer Acquisitions and Beyond: Policy Effects on Innovation Strategies. *International Economic Review*, **65** (2), 591–622.

LIU, O. (2022). The Dynamics of Firm Size Inequality: The Role of Acquisition and Innovation. *Mimeo*.

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. (2025). Product Differentiation and Oligopoly: A Network Approach. *American Economic Review*, **115** (4), 1170–1225.

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.

SEDLÁČEK, P. and STERK, V. (2017). The Growth Potential of Startups over the Business Cycle. *American Economic Review*, **107** (10), 3182–3210.

SERU, A. (2014). Firm Boundaries Matter: Evidence from Conglomerates and R&D Activity. *Journal of Financial Economics*, **111** (2), 381 – 405.

STERK, V., SEDLÁČEK, P. and PUGSLEY, B. (2021). The Nature of Firm Growth. *American Economic Review*, **111** (2), 547–79.

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

WEI, S.-J., XU, J., YIN, G. and ZHANG, X. (2023). Market for Patents, Monopoly, and Misallocation. *Mimeo*.

WEISS, J. (2022). Market Concentration, Growth, and Acquisitions. *Mimeo*.

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 Model Appendix

A.1 Timing of events for startup ideas

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

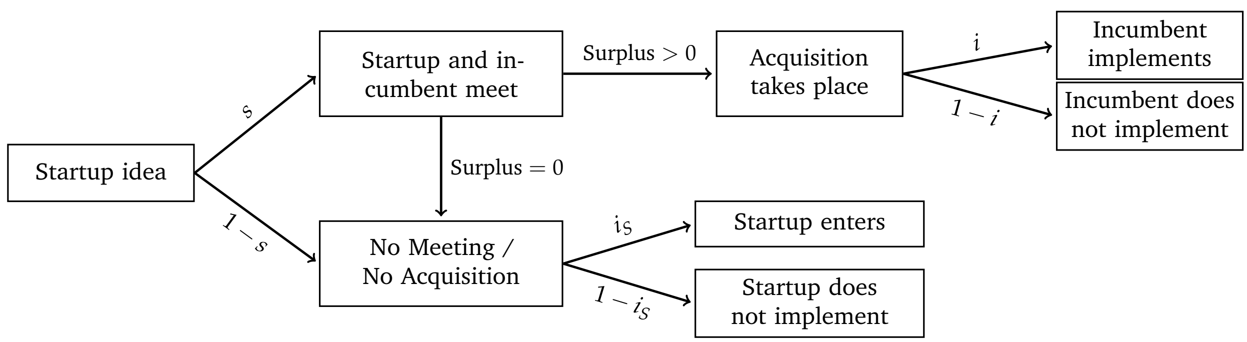

A.2 Value functions

The problem of the incumbent 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 its products. The value function of this firm solves the following HJB equation:[37]

$$r_t V_t(\mathbf{n}, n) = \max_{\substack{\{z_j, s_{R,j}\}_{j=1}^{n};\\ s_U;\, x}} \left\{ \sum_{j=1}^{n} \left[\underbrace{\omega_j \left(1 - \lambda^{-a_j}\right) Y_t}_{\text{Profits on product } j} - \underbrace{\xi_I^{\text{int}} z_j^{\psi} Y_t}_{\substack{\text{Internal innov.}\\ \text{costs on product } j}} - \underbrace{\chi\left(s_{R,j}^{\varphi} + s_U^{\varphi}\right) Y_t}_{\substack{\text{Costs of searching for}\\ \text{related and unrelated startups}}} \right.\right.$$

$$\left.\left. + \underbrace{z_j \left(V_t\left(\mathbf{n} \backslash \{(\omega_j, a_j)\} \cup \{(\omega_j, a_j + 1)\}, n\right) - V_t(\mathbf{n}, n) \right)}_{\text{Incumbent innovates on product } j} + \underbrace{x_S V_{R,jt}(\mathbf{n}, n)}_{\substack{\text{A startup has}\\ \text{an idea on product } j}}\right.\right.$$

[37] In this equation, $\cup$ and $\backslash$ are multiset union and difference operators. That is, 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 $\{a, b, b\} \backslash \{b\} = \{a\}$.

$$+ \underbrace{\widetilde{x}\left(V_t\Big(\mathbf{n}\backslash\{(\omega_j, a_j)\}, n-1\Big) - V_t(\mathbf{n}, n) \right)}_{\text{Loss of product } j \text{ to another incumbent}} \Bigg] + \underbrace{n x_S \gamma \int_0^1 V_{U,j't}(\mathbf{n}, n) dj'}_{\substack{\text{Startup with an idea on an unrelated} \\ \text{product is matched to the incumbent}}}$$

$$- \underbrace{\xi_I^{\text{ext}} n x^{\psi} Y_t}_{\substack{\text{External innov.} \\ \text{costs}}} + \underbrace{nx\left(\int_0^1 V_t\Big(\mathbf{n} \cup \{(\omega_{j'}, 1)\}, n+1\Big) \mathrm{d}j' - V_t(\mathbf{n}, n) \right)}_{\text{Incumbent externally innovates}} \Bigg\} + \dot{V}_t(\mathbf{n}, n). \quad \text{(A.1)}$$

In this equation, $V_{R,jt}(\mathbf{n}, n)$ is the change in the value of the incumbent when a startup has an idea on one of its products j, and $V_{U,j't}(\mathbf{n}, n)$ is the change in the value of the incumbent when it is matched to an unrelated startup with an idea on a product j'.

The value of a product We guess that there exists a product-level value function $\widehat{V}_t(\omega_j, a_j)$ that holds

$$V_t(\mathbf{n}, n) = \sum_{j=1}^{n} \widehat{V}_t(\omega_j, a_j).$$

and that the changes in value $V_{R,jt}(\mathbf{n}, n)$ and $V_{U,j't}(\mathbf{n}, n)$ only depend on the characteristics of products j and j'. This yields a product-level HJB equation:

$$\begin{aligned} r_t \widehat{V}_t(\omega_j, a_j) - \dot{\widehat{V}}_t(\omega_j, a_j) = \max_{z_j, s_{R,j}, s_U, x} \Bigg\{ & \omega_j \left(1 - \lambda^{-a_j}\right) Y_t - \xi_I^{\text{int}} z_j^{\psi} Y_t - \chi\Big(s_{R,j}^{\varphi} + s_U^{\varphi}\Big) Y_t \\ & + z_j \left(\widehat{V}_t(\omega_j, a_j + 1) - \widehat{V}_t(\omega_j, a_j) \right) + x_S \widehat{V}_{R,t}(\omega_j, a_j) - \widetilde{x} \widehat{V}_t(\omega_j, a_j) \\ & + x_S \gamma \int_0^1 \widehat{V}_{U,t}(\omega_{j'}) dj' - \xi_I^{\text{ext}} x^{\psi} Y_t + x \int_0^1 \widehat{V}_t(\omega_{j'}, 1) dj' \Bigg\}, \end{aligned} \quad \text{(A.2)}$$

where $\widehat{V}_{R,t}(\omega_j, a_j)$ and $\widehat{V}_{U,t}(\omega_{j'})$ are now the product-level change in incumbent's value from the acquisition of an idea on a related product and on an unrelated product.

As we will show later, these term depends on the surplus from acquiring the startup. The surplus is given by the joint post-acquisition value of both firms, net of their outside options. These are equal to:

$$\widehat{B}_{U,t}(\omega_j) = \max_{i_{U,j}} \left\{ i_{U,j} \widehat{V}_t(\omega_j, 1) - \kappa_I i_{U,j}^{\theta} Y_t \right\} - \max_{i_{S,j}} \left\{ i_{S,j} \widehat{V}_t(\omega_j, 1) - \kappa_S i_{S,j}^{\theta} Y_t \right\} - \tau_U p_{U,jt}^{A} Y_t \quad \text{(A.3)}$$

$$\widehat{B}_{R,t}(\omega_j, a_j) = \max_{i_{R,j}} \left\{ i_{R,j} \left(\widehat{V}_t(\omega_j, a_j + 1) - \widehat{V}_t(\omega_j, a_j) \right) - \kappa_I i_{R,j}^{\theta} Y_t \right\} - \tau_R p_{R,jt}^{A} Y_t \quad \text{(A.4)}$$

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

Taking the first-order conditions of the three sub-problems, we obtain:

$$i_{R,jt} = \left(\frac{\left(\widehat{V}_t(\omega_j, a_j+1) - \widehat{V}_t(\omega_j, a_j)\right)/Y_t}{\kappa_I \theta} \right)^{\frac{1}{\theta-1}}, \quad i_{U,jt} = \left(\frac{\widehat{V}_t(\omega_j, 1)/Y_t}{\kappa_I \theta} \right)^{\frac{1}{\theta-1}}, \quad i_{S,jt} = \left(\frac{\widehat{V}_t(\omega_j, 1)/Y_t}{\kappa_S \theta} \right)^{\frac{1}{\theta-1}}. \tag{A.5}$$

Plugging these back into equations (A.3)-(A.4), we can write the surpluses as follows:

$$\widehat{B}_{U,t}(\omega_j) = \kappa_I(\theta - 1) i_{U,jt}^{\theta} Y_t - \tau_U p_{U,jt}^A Y_t - \kappa_S(\theta - 1) i_{S,jt}^{\theta} Y_t, \tag{A.6}$$

$$\widehat{B}_{R,t}(\omega_j, a_j) = \kappa_I(\theta - 1) i_{R,jt}^{\theta} Y_t - \tau_R p_{R,jt}^A Y_t + i_{S,jt}\widehat{V}_t(\omega_j, a_j) - \kappa_S(\theta - 1) i_{S,jt}^{\theta} Y_t. \tag{A.7}$$

Nash bargaining solutions At this stage, we can determine the acquisition prices for related and unrelated startups. For unrelated acquisitions, the bargaining problem is:

$$\max_{p_{U,jt}^A} \left(\kappa_I(\theta - 1) i_{U,jt}^{\theta} - (1 + \tau_U) p_{U,jt}^A \right)^{\alpha} \left(p_{U,jt}^A - \kappa_S(\theta - 1) i_{S,jt}^{\theta} \right)^{1-\alpha} Y_t$$
$$\text{subject to } p_{U,jt}^A \geq \kappa_S(\theta - 1) i_{S,jt}^{\theta}$$

Solving for the optimal price, we obtain

$$p_{U,jt}^A = \left((1 - \alpha) \left(\frac{\kappa_I(\theta - 1) i_{U,jt}^{\theta}}{1 + \tau_U} \right) + \alpha \left(\kappa_S(\theta - 1) i_{S,jt}^{\theta} \right) \right) Y_t$$

Plugging the price back into equation (A.6) gives

$$\widehat{B}_{U,t}(\omega_j) = (1 + \alpha\tau_U) \left[\frac{\kappa_I(\theta - 1) i_{U,jt}^{\theta}}{1 + \tau_U} - \kappa_S(\theta - 1) i_{S,jt}^{\theta} \right] Y_t,$$

which is the expression for the surplus shown in the main text (conditional on a normalization with respect to aggregate output). Combining these results, the change in the value of the firm after being matched to an unrelated startup can be written as:

$$\widehat{V}_{U,t}(\omega_j) \equiv s_U \widetilde{\alpha}_U \max\left(\widehat{B}_{U,t}(\omega_j), 0 \right), \tag{A.8}$$

where $\widetilde{\alpha}_U$ is defined as in the main text.

For related acquisitions, the bargaining problem is

$$\max_{p^A_{R,jt}} \left(\kappa_I(\theta-1)i^{\theta}_{R,jt} - (1+\tau_R)p^A_{R,jt} + i_{S,jt}\widehat{V}_t(a_j,\omega_j)\right)^{\alpha}\left(p^A_{R,jt} - \kappa_S(\theta-1)i^{\theta}_{S,jt}\right)^{1-\alpha} Y_t$$
$$\text{subject to } p^A_{R,jt} \geq \kappa_S(\theta-1)i^{\theta}_{S,jt}$$

Similarly to the case of unrelated acquisitions, the price equals

$$p^A_{R,jt} = \kappa_S(\theta-1)i^{\theta}_{S,jt}Y_t + (1-\widetilde{\alpha}_R)\widehat{B}_{R,t}(\omega_j,a_j),$$

where $\widetilde{\alpha}_R$ is defined in the main text. Thus, the surplus can be written as

$$\widehat{B}_{R,t}(\omega_j,a_j) = (1+\alpha\tau_R)\left[(\theta-1)\left(\frac{\kappa_I i^{\theta}_{R,jt}}{1+\tau_R} - \kappa_S i^{\theta}_{S,jt}\right)Y_t + \frac{i_{S,jt}\widehat{V}_t(\omega_j,a_j)}{1+\tau_R}\right]$$

and, finally, the expected change in the value of the product for the incumbent is

$$\widehat{V}_{R,t}(\omega_j,a_j) \equiv s_{R,jt}\widetilde{\alpha}_R(1-\gamma)\max\left(\widehat{B}_{R,t}(\omega_j),0\right) - \left(\gamma s_U i_{U,jt} + (1-\gamma s_U)i_{S,jt}\right)\widehat{V}_t(\omega_j,a_j)$$

Normalized product value Finally, we guess-and-verify that there exist time-invariant functions $v : \Omega\times\mathbb{N}\to\mathbb{R}$, $b_R : \Omega\times\mathbb{N}\to\mathbb{R}$ and $b_U : \Omega\to\mathbb{R}$ such that $\widehat{V}_t(\omega,a) = v(\omega,a)Y_t$, $\widehat{B}_{R,t}(\omega,a) = b_R(\omega,a)Y_t$ and $\widehat{B}_{U,t}(\omega) = b_U(\omega)Y_t$. This guess implies that $\dot{\widehat{V}}_t(\omega_j,a_j) = v(\omega_j,a_j)gY_t$, where $g = \dot{Y}_t/Y_t$ is the growth rate of aggregate output. Under this guess, the optimal implementation rates from equation (A.5) now read

$$i_{R,jt} = \left(\frac{v(\omega,a+1)-v(\omega,a)}{\kappa_I\theta}\right)^{\frac{1}{\theta-1}}, \quad i_{U,jt} = \left(\frac{v(\omega,1)}{\kappa_I\theta}\right)^{\frac{1}{\theta-1}}, \quad i_{S,jt} = \left(\frac{v(\omega,1)}{\kappa_S\theta}\right)^{\frac{1}{\theta-1}}. \tag{A.9}$$

so we can write $i_{R,jt} = i_R(\omega_j,a_j)$, $i_{U,jt} = i_R(\omega_j)$ and $i_{S,jt} = i_S(\omega_j)$. Note that all products j with the same (ω_j,a_j) have the same value, which allows us to drop the j subscript everywhere. Then, plugging our guess into equation (A.2), using $r = \rho+g$ by the Euler equation, and dividing by Y_t, we obtain:

$$\rho v(\omega,a) = \max_{z,s_R,s_U,x}\Bigg\{\omega(1-\lambda^{-a}) - \zeta^{\text{int}}_I z^{\psi} - \chi\left(s^{\varphi}_R + s^{\varphi}_U\right) + z\Big(v(\omega,a+1)-v(\omega,a)\Big)$$
$$+ x_S\Big(s_R\widetilde{\alpha}_R(1-\gamma)\max\left(b_R(\omega,a),0\right) - d(\omega)v(\omega,a)\Big) - \widetilde{x}v(\omega,a)$$

$$+ x_S s_U \widetilde{\alpha}_U \gamma \sum_{\omega' \in \Omega} \phi(\omega') \max\Big(b_U(\omega'), 0\Big) - \xi_I^{\text{ext}} x^\psi + x \sum_{\omega' \in \Omega} \phi(\omega') v(\omega', 1) \Bigg\}, \tag{A.10}$$

where $d(\omega) \equiv \gamma s_U i_U(\omega) + (1 - \gamma s_U) i_S(\omega)$. Taking first-order conditions yields equations (8), (9), (18) and (25) in the main text. Plugging back these results back into equation (A.10), and using the equilibrium condition $x = \widetilde{x}$, we get equation (26).

A.3 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, and the mass of products with a given spending share, $\phi(\omega)$, is an exogenous parameter.

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

$$\begin{aligned} &\Big(z(\omega, a-1) + x_S (1-\gamma) s_R(\omega, a-1) i_R(\omega, a-1)\Big) m(\omega, a-1) = \\ &\qquad \Big(z(\omega, a) + x_S i(\omega, a) + x\Big) m(\omega, a). \end{aligned} \tag{A.11}$$

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 by themselves or acquired from a related startup. Any innovation triggers an outflow.

Moreover, the distribution must hold

$$\sum_{a=1}^{+\infty} m(\omega, a) = \phi(\omega). \tag{A.12}$$

Equations (A.11) and (A.12) define a linear system that can be solved for the distribution m once we know all innovation rates and implementation and noticing probabilities.

A.4 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

$$I_t^{\text{Inc}} = Y_t \Bigg[\xi_I^{\text{ext}} x^\psi + \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) \Big(\xi_I^{\text{int}} \big(z(\omega, a)\big)^\psi + \dots$$

$$\cdots + x_S\Big[(1-\gamma)s_R(\omega,a)\kappa_I\big(i_R(\omega,a)\big)^\theta + \gamma s_U\kappa_I\big(i_U(\omega)\big)^\theta\Big]\Big)\Bigg], \quad \text{(A.13)}$$

$$I_t^{\text{Startup}} = Y_t\left[x_S\left(\xi_S + \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m(\omega,a)\Big(\gamma(1-s_U) + (1-\gamma)(1-s_R(\omega,a))\Big)\kappa_S\big(i_S(\omega)\big)^\theta\right)\right], \quad \text{(A.14)}$$

$$S_t = Y_t\left[\chi(s_U)^\varphi + \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m(\omega,a)\chi\big(s_R(\omega,a)\big)^\varphi\right]. \quad \text{(A.15)}$$

All costs scale linearly in Y_t. This shows that consumption is always proportional to output, and thus grows at the same rate.

Growth rate By equation (30), output, consumption and wages grow at the same rate as aggregate productivity Q_t on the BGP. Using the definition of aggregate productivity, we can write:

$$Q_t = \exp\left(\int_0^1 \omega_j \ln(\omega_j)\,\mathrm{d}j\right)\underbrace{\exp\left(\int_0^1 \omega_j \ln(q_{jt})\,\mathrm{d}j\right)}_{\equiv \mathcal{Q}_t}. \quad \text{(A.16)}$$

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

$$\ln(\mathcal{Q}_{t+dt}) = \int_0^1 \omega_j \ln(q_{j,t+dt})\mathrm{d}j. \quad \text{(A.17)}$$

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

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

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 $\Big(x + z(\omega,a) + x_S i(\omega,a)\Big)dt$ 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,

$$\frac{\ln(\mathcal{Q}_{t+dt}/\mathcal{Q}_t)}{dt} = \ln(\lambda)\sum_{\omega\in\Omega}\omega\left[\sum_{a=1}^{+\infty} m(\omega,a)\Big(x + z(\omega,a) + x_S i(\omega,a)\Big)\right]. \quad \text{(A.18)}$$

Taking the limit as $dt \to 0$ and rearranging terms gives equation (31).

A.5 BGP solution algorithm

To solve for the BGP equilibrium, we define a function h which gives the error in the free entry condition, equation (28), as a function of the startup rate. That is,

$$h(x_S) = \xi_S - \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m(\omega,a)\Bigg[\kappa_S\,(\theta-1)\,(i_S(\omega))^{\theta} + \dots$$
$$\dots + (1-\gamma)\,s_R(\omega,a)\,(1-\widetilde{\alpha}_R)\,b_R\,(\omega,a) + \gamma s_U\,(1-\widetilde{\alpha}_U)\,b_U(\omega)\Bigg].$$

The equilibrium startup rate is the zero of this error function, and we solve for it using the `fzero` routine in Matlab. To compute the value of the error function for a given startup rate x_S, we need to solve the incumbent's dynamic problem and determine the invariant distribution over spending shares and technology gaps m. To do so, we proceed as follows.

1. Solve for the product-level value function through a value function iteration on the HJB equation (7):

   (a) Guess a value function, $v : \Omega \times \mathbb{A} \to \mathbb{R}$, where $\mathbb{A} \equiv \{1, 2, \dots, a_{\max}\}$ and $a_{\max}$ is a sufficiently large integer.[38]

   (b) Using equations (8), (9), (11), (13), (15), (18), (20), (22) and (25), deduce 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.

   (c) Using equation (7), compute the implied value of the value function v, and compare it to the current guess. If the implied value and the guess do not coincide, update the guess according to the rule

   $$v^{\text{new}} = 0.05 v^{\text{implied}} + 0.95 v^{\text{old}}.$$

2. Using the innovation rates obtained in step 1 and the assumed value for the startup rate x_S, compute the joint distribution of products across spending shares and technology gaps m as described in section A.3.

Note that the algorithm is independent of aggregate outcomes such as the growth rate g. Thus, we can solve for these outcomes after computing firm-level decisions, using the equations in Section 2.2.7. This block structure greatly simplifies the solution of the model.

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

A.6 Transition between BGPs

To characterize the transition from one BGP to another, we use a discrete-time approximation of our continuous-time model. Section A.6.1 describes this approximation, and Section A.6.2 lays out the algorithm used to compute the transition path.

A.6.1 A discrete-time approximation

Outside of the BGP, the product-level HJB equation holds

$$
\begin{aligned}
(r_t - g_t)\, v_t(\omega,a) - \dot{v}_t(\omega,a) = \max_{z_t, s_{R,t}, s_{U,t}, x_t} \Bigg\{ & \omega\left(1-\lambda^{-a}\right) - \zeta_I^{\text{int}} z_t^{\psi} - \chi\left(s_{R,t}^{\varphi} + s_{U,t}^{\varphi}\right) - \zeta_I^{\text{ext}} x_t^{\psi} \\
& + z_t\Big(v_t(\omega,a+1) - v_t(\omega,a)\Big) + x_t \sum_{\omega'\in\Omega} \phi(\omega') v_t(\omega',1) - \tilde{x}_t v_t(\omega,a) \\
& + x_{S,t} v_{R,t}(s_{R,t},\omega,a) + x_{S,t}\gamma v_{U,t}\left(s_{U,t}\right) \Bigg\}.
\end{aligned}
\tag{A.19}
$$

We now construct a discrete-time approximation to this continuous-time equation by dividing each year into short periods of length Δ. Approximating a Poisson rate z_t for an event by the discrete probability Δz_t for this event happening within a period, we get

$$
\begin{aligned}
v_t(\omega,a) = \max_{z_t, s_{R,t}, s_{U,t}, x_t} \Bigg\{ & \Delta\left(\omega\left(1-\lambda^{-a}\right) - \zeta_I^{\text{int}} z_t^{\psi} - \chi\left(s_{R,t}^{\varphi} + s_{U,t}^{\varphi}\right) - \zeta_I^{\text{ext}} x_t^{\psi}\right) \\
& + \exp(-\Delta \tilde{r}_{t+\Delta})\Bigg[\Delta z_t\, v_{t+\Delta}(\omega,a+1) + \Delta x_t \sum_{\omega'\in\Omega} \phi(\omega') v_{t+\Delta}(\omega',1) \\
& + \Big(1-\Delta\left(z_t + \tilde{x}_t\right)\Big) v_{t+\Delta}(\omega,a) + \Delta x_{S,t} v_{R,t+\Delta}(s_{R,t},\omega,a) + \Delta x_{S,t}\gamma v_{U,t+\Delta}\left(s_{U,t}\right)\Bigg]\Bigg\},
\end{aligned}
$$

where $\tilde{r}_{t+\Delta} \equiv r_{t+\Delta} - g_{t+\Delta}$ is the effective discount factor.

In continuous time, firms decide on innovation, search and implementation within the same instant. In discrete time, we need to make further assumptions on timing.[39] We assume that at time t, incumbent firms decide on the innovation rates z_t and x_t and on the search efforts $s_{R,t}$ and $s_{U,t}$. Likewise, the choices of startups imply a startup rate $x_{S,t}$. The resource costs for all of these expenses are paid at time t.

At time $t+\Delta$, the outcome of these investments is realized. That is, incumbents

[39] As time periods are very short in our numerical implementation, these assumptions are irrelevant for our results.

learn whether they have made an innovation, and whether they have met a (related or unrelated) startup. Likewise, startups know whether they made an innovation and whether they have been noticed by an incumbent. At that point, incumbents and startups decide on acquisitions and on the implementation probabilities $i_{S,t+\Delta}$, $i_{U,t+\Delta}$ and $i_{R,t+\Delta}$. The implementation lotteries are then realized immediately, within period.

With these assumptions, the innovation policy functions of incumbents hold

$$z_t(\omega,a) = \left(\exp(-\Delta\widetilde{r}_{t+\Delta})\frac{v_{t+\Delta}(\omega,a+1)-v_{t+\Delta}(\omega,a)}{\zeta_I^{\text{int}}\psi}\right)^{\frac{1}{\psi-1}}$$

$$x_t = \left(\exp(-\Delta\widetilde{r}_{t+\Delta})\frac{\sum_{\omega'\in\Omega}\phi(\omega')v_{t+\Delta}(\omega',1)}{\zeta_I^{\text{ext}}\psi}\right)^{\frac{1}{\psi-1}}.$$

These are essentially the same expressions as in the main text, expect that there is now discounting to account for the fact that costs and payoffs are realized in different periods.

The implementation probabilities for ideas, in turn, hold

$$i_{S,t+\Delta}(\omega') = \left(\frac{v_{t+\Delta}(\omega',1)}{\kappa_S\theta}\right)^{\frac{1}{\theta-1}}.$$

$$i_{U,t+\Delta}(\omega') = \left(\frac{v_{t+\Delta}(\omega',1)}{\kappa_I\theta}\right)^{\frac{1}{\theta-1}}.$$

$$i_{R,t+\Delta}(\omega,a) = \left(\frac{v_{t+\Delta}(\omega,a+1)-v_{t+\Delta}(\omega,a)}{\kappa_I\theta}\right)^{\frac{1}{\theta-1}}.$$

The surplus from an unrelated acquisition, the associated continuation value and the search effort hold

$$b_{U,t+\Delta}(\omega') = (1+\alpha\tau_U)(\theta-1)\left(\frac{\kappa_I}{1+\tau_U}\left(i_{U,t+\Delta}(\omega')\right)^{\theta}-\kappa_S\left(i_{S,t+\Delta}(\omega')\right)^{\theta}\right),$$

$$v_{U,t+\Delta}(s_{U,t}) = s_{U,t}\widetilde{\alpha}_U\sum_{\omega'\in\Omega}\phi(\omega')\max\left(0,b_{U,t+\Delta}(\omega')\right).$$

$$s_{U,t} = \left[\exp(-\Delta\widetilde{r}_{t+\Delta})\frac{\widetilde{\alpha}_U\gamma x_{S,t}}{\chi\varphi}\left(\sum_{\omega'\in\Omega}\phi(\omega')\max\left(0,b_{U,t+\Delta}(\omega')\right)\right)\right]^{\frac{1}{\varphi-1}}.$$

The surplus from a related acquisition and the associated continuation value hold

$$b_{R,t+\Delta}(\omega,a) = (1+\alpha\tau_R)\left(\frac{i_{S,t+\Delta}(\omega)v_{t+\Delta}(\omega,a)}{1+\tau_R} + (\theta-1)\left(\frac{\kappa_I\left(i_{R,t+\Delta}(\omega,a)\right)^\theta}{1+\tau_R} - \kappa_S\left(i_{S,t+\Delta}(\omega)\right)^\theta\right)\right),$$

$$\begin{aligned} v_{R,t+\Delta}(s_{R,t},\omega,a) = & -\Big(\gamma s_{U,t} i_{U,t+\Delta}(\omega) + (1-\gamma s_{U,t})\, i_{S,t+\Delta}(\omega)\Big) v_{t+\Delta}(\omega,a) \\ & + s_{R,t}\widetilde{\alpha}_R\,(1-\gamma)\max\Big(0, b_{R,t+\Delta}(\omega,a)\Big). \end{aligned}$$

The search effort for related startups is

$$s_{R,t}(\omega,a) = \left[\exp(-\Delta\widetilde{r}_{t+\Delta})\frac{\widetilde{\alpha}_R\,(1-\gamma)\,x_{S,t}}{\chi\varphi}\max\left(0, b_{R,t+\Delta}(\omega,a)\right)\right]^{\frac{1}{\varphi-1}}.$$

This fully characterizes the decision problem of an incumbent firm. In turn, the expected value of a startup that pays the fixed cost of innovation at time t is given by

$$\begin{aligned} v_{S,t} = \; & \exp(-\Delta\widetilde{r}_{t+\Delta})\left(\sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m_t(\omega,a)\Bigg[\kappa_S\,(\theta-1)\,\left(i_{S,t+\Delta}(\omega)\right)^\theta\right. \\ & \left. + (1-\widetilde{\alpha}_R)\,(1-\gamma)\,s_{R,t}(\omega,a)b_{R,t+\Delta}\,(\omega,a) + (1-\widetilde{\alpha}_U)\,\gamma s_{U,t} b_{U,t+\Delta}(\omega)\Bigg]\right). \end{aligned}$$

With these definitions, we are now ready to describe the algorithm for computing outcomes during the transition from one BGP to another.

A.6.2 Transitional dynamics algorithm

We assume that the economy is initially on a baseline no-policy BGP. At time $t = 0$, it is hit by a permanent and unexpected shock which changes acquisition taxes to $\left(\tau_R^{\text{New}}, \tau_U^{\text{New}}\right)$. This shock will make the economy transition to a new BGP.

To characterize the transition period, we make two numerical approximations. First, we assume that the transition is completed (i.e., that the economy is very close to its new BGP) after a period of time of $T = 100$ years. Second, we consider the discrete-time version of the model described in the previous section, with a period length of $\Delta = 0.0005$. To solve the model, we use the following algorithm:

1. **Start of the outer loop.** At the first iteration, make an initial guess for the path of the effective discount rate, $\widetilde{\mathbf{r}}^{(0)} = \left\{\left(r_t^{(0)} - g_{Y,t}^{(0)}\right)\right\}_{t\in\{\Delta,2\Delta,\ldots,T\}}$, where $g_{Y,t}$ denotes the growth rate of output between period $t-\Delta$ and t, and r_t the interest rate between

period $t-\Delta$ and t. In practice, we set the initial guess for the effective discount rate to ρ, its value on any BGP.

2. **Start of the inner loop.** Guess a path for the startup rate $\mathbf{x}_S^{(0)} = \left\{ x_{S,t}^{(0)} \right\}_{t\in\{0,\Delta,2\Delta,\ldots,T-\Delta\}}$. For the first iteration of the outer loop, this initial guess is arbitrary. For further iterations of the outer loop, the initial guess is the path obtained at the end of the last run of the inner loop.

3. Given the guesses for the effective discount rates and startup rates, solve backwards for all firm decisions. To do so, assume that $v_T^{(k)} = v^{\text{New}}$, where v^{New} is the value function on the new BGP. Then, starting from $t = T-\Delta$ and going until $t = 0$ in steps of size Δ:

   (a) Given the value function $v_{t+\Delta}^{(k)}$, solve for innovation policy functions $z_t^{(k)}$ and $x_t^{(k)}$, implementation policy functions $i_{S,t+\Delta}^{(k)}$, $i_{U,t+\Delta}^{(k)}$ and $i_{R,t+\Delta}^{(k)}$, and search effort policy functions $s_{R,t}^{(k)}$ and $s_{U,t}^{(k)}$, using the equations derived in Section A.6.1.

   (b) Use the results from (a) to compute the value function $v_t^{(k)}$. Similarly as in the main text, we can show that this value function holds:

$$
\begin{aligned}
v_t^{(k)}(\omega,a) = \Delta\Bigg[&\pi(\omega,a) + (\psi-1)\left(\zeta_I^{\text{int}}\left(z_t^{(k)}(\omega,a)\right)^{\psi} + \zeta_I^{\text{ext}}\left(x_t^{(k)}\right)^{\psi}\right) \\
&+\chi(\varphi-1)\left(\left(s_{R,t}^{(k)}(\omega,a)\right)^{\varphi} + \left(s_{U,t}^{(k)}\right)^{\varphi}\right)\Bigg] + \exp\left(-\Delta\tilde{r}_{t+\Delta}^{(k)}\right) \\
&\cdot\left[\left(1-\Delta\left(x_t^{(k)} + x_{S,t}^{k}\left(\gamma s_{U,t}^{(k)} i_{U,t+\Delta}^{(k)}(\omega) + \left(1-\gamma s_{U,t}^{(k)}\right) i_{S,t+\Delta}^{(k)}(\omega)\right)\right)\right) v_{t+\Delta}^{(k)}\right].
\end{aligned}
$$

4. Starting from the initial distribution of technology gaps $m_0(\omega,a) = m(\omega,a)$, and using the firm policy functions computed in step 2, compute the distribution of technology gaps for any period between $t=\Delta$ and $t=T$. To do this, we use the discretized flow equations given below. For every $a>1$,

$$
\begin{aligned}
m_{t+\Delta}^{(k)}(\omega,a) = &\left(1-\Delta\left(x_t^{(k)} + z_t^{(k)}(\omega,a) + z_{S,t+\Delta}^{(k)}(\omega,a)\right)\right) m_t^{(k)}(\omega,a) \\
&+\Delta\left(z_t^{(k)}(\omega,a-1) + x_{S,t}^{k}(1-\gamma)\, s_{R,t}^{(k)}(\omega,a-1) i_{R,t+\Delta}^{(k)}(\omega,a-1)\right) m_t^{(k)}(\omega,a-1).
\end{aligned}
$$

Finally, for $a = 1$,

$$m_{t+\Delta}^{(k)}(\omega, 1) = \phi(\omega) - \sum_{a=2}^{+\infty} m_{t+\Delta}^{(k)}(\omega, a).$$

5. Compute the expected value of a startup idea for every period t, $v_{S,t}$, as defined in Section A.6.1. If for every t, this expected value is close enough to the startup creation cost ξ_S, then the inner loop of the algorithm has converged, and we proceed to step 6. Otherwise, update the guess for $\mathbf{x}_S$ according to

$$x_{S,t}^{(k+1)} = \left(\frac{v_{S,t}^{(k)}}{\xi_S}\right)^{\vartheta} x_{S,t}^{(k)},$$

for all $t \in \{0, \Delta, ..., T-\Delta\}$, where $\vartheta \in (0,1)$ is a dampening parameter, and go back to Step 3.

6. Compute the implied aggregate output and aggregate consumption levels, for each $t \in \{0, \Delta, ..., T-\Delta\}$.

   (a) For aggregate output, use equation (30):

   $$\mathbf{Y}_t^{(k)} = \mathcal{M}_t^{(k)} \left(\mathcal{W}_t^{(k)}\right)^{-1} Q_t^{(k)} L,$$

   where $\mathcal{M}_t^{(k)} = \left(\sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m_t^{(k)}(\omega, a)\omega\lambda^{-a}\right)^{-1}$, $\mathcal{W}_t^{(k)} = \lambda^{\sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m_t^{(k)}(\omega,a)a\omega}$, and $Q_t^{(k)}$ follows the law of motion

$$\ln(Q_{t+\Delta}^{(k)}) = \Delta\left(\ln(\lambda)\sum_{\omega\in\Omega}\omega\left[\sum_{a=1}^{+\infty} m_t^{(k)}(\omega, a)\left(x_t^{(k)} + z_t^{(k)}(\omega, a) + x_{S,t}^{(k)} i_t^{(k)}(\omega, a)\right)\right]\right) + \ln(Q_t^{(k)})$$

   where we normalize $Q_0^{(k)} = 1$.

   (b) For aggregate consumption, use $C_t^{(k)} = Y_t^{(k)}\left(1 - \frac{I_t^{\text{Inc},(k)}}{Y_t^{(k)}} - \frac{I_t^{\text{Startup},(k)}}{Y_t^{(k)}} - \frac{S_t^{(k)}}{Y_t^{(k)}}\right)$, where

$$\frac{I_t^{\text{Inc},(k)}}{Y_t^{(k)}} = \xi_I^{\text{ext}}(x_t^{(k)})^{\psi} + \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m_t^{(k)}(\omega, a)\xi_I^{\text{int}}\left(z_t^{(k)}(\omega, a)\right)^{\psi} + \dots$$

$$+ x_{S,t-\Delta}^{(k)}\sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m_{t-\Delta}^{(k)}(\omega, a)\left[(1-\gamma)s_{R,t-\Delta}^{(k)}(\omega, a)\kappa_I\left(i_{R,t}^{(k)}(\omega, a)\right)^{\theta} + \gamma s_{U,t-\Delta}^{(k)}\kappa_I\left(i_{U,t}^{(k)}(\omega)\right)^{\theta}\right],$$

$$\frac{I_t^{\text{Startup},(k)}}{Y_t^{(k)}} = x_{S,t}^{(k)}\xi_S + x_{S,t-\Delta}^{(k)}\left(\sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m_{t-\Delta}^{(k)}(\omega, a)\right.$$

$$\cdot \left(\gamma(1-s_{U,t-\Delta}^{(k)}) + (1-\gamma)(1-s_{R,t-\Delta}^{(k)}(\omega,a))\right)\kappa_S\left(i_{S,t}^{(k)}(\omega)\right)^{\theta}\Big),$$

$$\frac{S_t^{(k)}}{Y_t^{(k)}} = \chi(s_{U,t}^{(k)})^{\varphi} + \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m_t^{(k)}(\omega,a)\chi\left(s_{R,t}^{(k)}(\omega,a)\right)^{\varphi}.$$

7. For each $t \in \mathcal{T}$, compute the implied growth rates of output and consumption, and get a new interest rate from the Euler equation:

$$g_{Y,t} = \frac{1}{\Delta}\left(\frac{Y_t^{(k)}}{Y_{t-\Delta}^{(k)}} - 1\right), \qquad \text{and} \qquad r_t = \frac{1}{\Delta}\left(\frac{C_t^{(k)}}{C_{t-\Delta}^{(k)}} - 1\right) + \rho.$$

8. Collect these results in a vector $\widetilde{\mathbf{r}} = \left\{\left(r_t - g_{Y,t}\right)\right\}$. Stop if $\|\widetilde{\mathbf{r}} - \widetilde{\mathbf{r}}^{(k)}\| < \varepsilon$, for some small tolerance $\varepsilon > 0$, where $\|\cdot\|$ denotes the sup-norm. Otherwise, define:

$$\widetilde{\mathbf{r}}^{(k+1)} = \iota\widetilde{\mathbf{r}} + (1-\iota)\widetilde{\mathbf{r}}^{(k)},$$

   where $\iota \in (0,1)$ is a dampening parameter, and go back to Step 2 with $[k] \leftarrow [k+1]$.

A.7 Decomposition formula

Here, we derive the decomposition for changes in the growth rate between BGPs in Section 5.1. We can write the BGP growth rate in equation (31) as

$$g = \ln(\lambda)\Big(\text{Incumbent own innovation} + x_S \cdot \text{Impl. rate of startup ideas}\Big), \tag{A.20}$$

where incumbent own innovation and the implementation rate of startup ideas are defined as in equation (37) in the main text. Taking the ratio of growth rates in the policy and no-policy BGPs, we immediately get equation (37), with

$$\text{Share}_{\text{inc}}^{A} \equiv \frac{\ln(\lambda)\cdot\text{Incumbent own innovation}^{\text{No Policy}}}{g^{\text{No Policy}}}. \tag{A.21}$$

B Data appendix

B.1 SDC Platinum

In SDC Platinum, we select all completed deals between 2000 and 2020 in which both target and acquirer are located in the United States.[40] We drop deals in which the ultimate parent firms of target or acquirer belong to the Finance, Insurance, Utilities, Government or Public Administration sectors. We also drop deals in which the target firm is a subsidiary (i.e., deals where the SDC firm identifier for the target firm is different from the one of their ultimate parent firm).

To identify deals that qualify as startup acquisitions, we impose two conditions:

1. At the time of acquisition, the target firm must be at most six years old.
2. The percentage of target firm shares acquired in the deal must exceed 50%.

To compute target age, we use the SDC variables "Date Announced" and "Target Founded Date". While SDC always provides the announcement date, the "Target Founded Date" variable has unfortunately many missing values. As a result, we a priori have to exclude all of these deals, even though many of them may be startup acquisitions.

Figure B.1: Number of startup acquisitions per year (2000-2020)

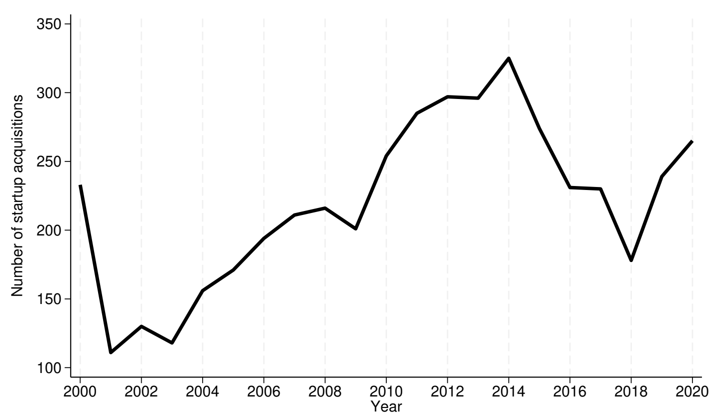

To alleviate this issue, we recover some missing founding dates from Ewens and Marx (2024), who provide founding dates for a large number of patenting firms. To do so, we match all target firms in our SDC sample to patenting firms in PatentsView, using the procedure described in Section B.2 below. Then, if a matched firm has no founding date

[40] SDC is at https://www.lseg.com/en/data-analytics/products/sdc-platinum-financial-securities.

in SDC Platinum, but does have a founding date in the Ewens and Marx dataset, we use the latter to fill out the missing SDC value.[41] This procedure adds roughly 600 startup acquisitions to the sample. Overall, our sample contains 4615 startup acquisition deals. Table 1 in the main text shows some summary statistics for this sample, and Figure B.1 plots the number of startup acquisitions over time.

B.2 Matching startups to patents

We aim to match all potential startup targets (i.e., firms which were founded at most six years before their acquisition, or firms with missing founding years) to their patents, if they hold any. As there is no crosswalk between SDC Platinum and PatentsView, we rely on firm names, the only common information across these two datasets.[42]

To carry out the match, we first use the name standardization routine `stnd_compname` developed by Wasi and Flaaen (2015) for Stata. Among other things, this routine removes commonly used endings for firm names, such as "CO", "CORP" or "CORPORATION". Then, we match standardized firm names in SDC to standardized firm names in PatentsView, keeping only exact matches (in which standardized firm names are identical in both datasets).[43]

B.3 Matching acquirers to Compustat

We download the Compustat Fundamentals Annual database through Wharton Research Data Services. We restrict the sample to US firms that do not belong to the Finance, Insurance, Utilities or Government sectors (as in our acquisition data), and consider the time period between 2000 and 2020. We also drop subsidiaries (i.e., observations with stock ownership codes (STKO) equal to 1 or 2), observations with missing employment or sales, as well as observations for firms with less than 10 employees, less than $10,000 of sales or negative R&D spending.

We use this data to match Compustat acquirers to their acquisitions of startups recorded in SDC Platinum. To do so, we rely on the work of Ewens *et al.* (2025), which provides a

[41] That is, for startups with founding dates in both datasets, we always use SDC Platinum. Note that Ewens and Marx report the founding year, while SDC reports the exact founding date. To bring the data into the same format, we set the founding date to June 30 for all observations in the Ewens and Marx dataset.

[42] PatentsView data can be accessed at https://patentsview.org/.

[43] In a robustness check, we use the `reclink2` routine, also by Wasi and Flaaen (2015), to identify further potential matches. Potential matches are cases in which standardized names are similar, and assigned a "similarity score" by the algorithm. We then manually check potential matches with high similarity scores. This allow us to obtain an even larger sample, but does not affect our empirical results.

mapping between SDC Platinum deal numbers and Compustat firm identifiers (GVKEY). We use the 2024 version of their bridge file, which covers our entire sample period.[44]

B.4 Matching acquired to control patents

As shown in Table 1 in the main text, our matched SDC-PatentsView dataset has 2127 acquired startup patents, belonging to 620 startups. We restrict this sample to patents that were filed at least one year before the acquisition. For each of these treated patents, we identify potential controls as patents that are identical to them for six matching variables:

1. The patent application year.
2. The patent's main technology class.
3. The number of inventors on the patent.
4. The US state in which inventors are located.
5. The founding year of the firm that filed the patent.[45]
6. A dummy for whether the patent's backward similarity score (as computed by Kelly *et al.*, 2021), a measure of patent novelty, is above or below the median.

When there are 10 patents or less that hold these criteria, we keep all of them. When there are more, we keep a random sample of 10 patents. We are able to match 778 patents to a control group. By construction, acquired and control patents have the exact same values for the six matching variables.

Table B.1: Comparing Acquired and Control Patents

	Control Patents		Acquired Patents			
	Mean	Obs	Mean	Obs	Difference	p-value
ln(Number of Claims)	2.86	4371	2.88	778	0.02	0.56
Citations Received after 1 Year	0.54	4371	0.50	778	-0.04	0.61
New Bigrams	10.06	4089	10.02	733	-0.05	0.35
Textual Novelty	15.63	4371	15.53	778	-0.11	0.47

Notes: This table compares acquired patents to control patents for three variables which are not used in the matching exercise (number of claims, citations received after one year, and new bigrams), as well as for the exact value of the backward similarity score of Kelly *et al.* (2021).

[44] The updated file is at `https://github.com/michaelewens/SDC-to-Compustat-Mapping`.
[45] For control patents, data on the founding year comes again from Ewens and Marx (2024).

In Table B.1, we show that there are also no statistically significant differences for other observable characteristics that were not used for matching. These include the number of claims in the patent,[46] the number of citations received one year after filing, the number of new bigrams in the patent text (from Arts *et al.*, 2021) and the numerical value of the backward similarity score (from Kelly *et al.*, 2021). For each of these variables, a *t*-test for the difference in means between acquired patents and their control group is not significant, supporting the claim that both sets of patents are similar across observables.

B.5 Additional empirical results

B.5.1 Robustness checks on the main result

In Table B.2, we show a number of robustness checks on our main differences-in-differences regression (i.e., on Table 2 in the main text). All columns in this robustness table include both year fixed effects and matched group fixed effects.

For comparison purposes, Column (1) replicates Column (4) of Table 2, showing our main finding: on average, acquired patents do not receive a boost in citations after acquisition relative to their control patents. Column (2) increases the time window considered for our regression to 7 years prior to the acquisition and 7 years after the acquisition (our baseline window used 5 years). Column (3) includes industry-year fixed effects instead of year fixed effects. Column (4) uses an OLS instead of a Poisson regression, specifying the dependent variable in levels, and Column (5) uses an OLS regression where the dependent variable is specified in logs. In all specifications, our results are unchanged, and there is no statistically significant effect of acquisitions on patent citations.

Finally, in Column (6), we consider a new dependent variable: patent renewals. Four years after a patent has been granted, the patent owner has the option of paying a renewal fee to keep patent protection for an additional four years. Patent renewals are an alternative proxy for the acquirer's willingness to implement the ideas of the acquired startup: making a payment to renew a patent is a strong financial commitment that indicates the acquirer's aim to continue producing the products derived from this idea, or to conduct follow-on research in the same product line. We retrieve information on patent renewals from the USPTO's Open Data Portal,[47] and use it to define a dummy variable which is equal to 1 for all the years following the decision of a patent owner to renew, and 0 otherwise. We then

[46] 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).

[47] Precisely, we download the "Patent Maintenance Fee Events" database (Product Code PTMNFEE2) on `https://data.uspto.gov/bulkdata/datasets`.

estimate the baseline differences-in-differences specification shown in equation (32) using this alternative dependent variable. Aligned with our baseline findings, Column (6) shows that there is no significant effect of acquisition on the renewal of a startup patent.

Table B.2: Robustness: main regression results

	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.717***	0.905***	0.779***	0.108	0.062***	4.672***
	(0.112)	(0.125)	(0.110)	(0.091)	(0.022)	(0.789)
Acquired	0.064	0.049	0.039	0.027	0.005	0.421
	(0.310)	(0.312)	(0.295)	(0.069)	(0.021)	(0.292)
Acquired × Post	0.080	0.104	0.107	0.138	0.050	-0.174
	(0.287)	(0.293)	(0.273)	(0.132)	(0.036)	(0.306)
Matched Group FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Baseline	Yes					
Longer Trend		Yes				
Industry-Year FE			Yes			
OLS levels				Yes		
OLS logs					Yes	
Patent Renewal						Yes
Observations	28,576	35,730	27,812	32,282	32,282	31,637

Notes: Unless otherwise specified, 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. Column (1) repeats our baseline result. Column (2) expands the window of years to 7 years before the acquisition until 7 years after acquisition. Column (3) includes industry-year fixed effects instead of year fixed effects. Column (4) is an OLS regression in levels. Column (5) is an OLS regression in which the dependent variable is $\ln(1 + \text{NumCites})$. Finally, Column (6) replaces the dependent variable by a dummy for patent renewals. Standard errors are clustered at the target firm level. * significant at 10%; ** significant at 5%; *** significant at 1%.

B.5.2 Startup acquisitions and acquirer size

Table B.3 illustrates the link between startup acquisitions and acquirer firm size for the sample of Compustat firms with positive R&D spending. Column (1) shows the results for the specification shown in the main text in equation (35), reproduced below for convenience:

$$\text{SAQ}^s_{f,t} = \alpha^s_t + \beta_S \ln\left(\text{Emp}^s_{f,t}\right) + \varepsilon^s_{f,t},$$

where $\text{SAQ}^s_{f,t}$ is the number of startups acquired by firm f of industry s in year t, $\text{Emp}^s_{f,t}$ is the number of employees of firm f in year t, and α^s_t is a set of industry-year fixed effects (with industries defined at the 4-digit NAICS level). We normalize the dependent variable to have mean 1. As Column (1) shows, there is a strong positive correlation between the number of startups acquired and acquirer firm size.

Table B.3: The link between startup acquisitions and firm size

	(1)	(2)	(3)	(4)	(5)	(6)
ln(Employment)	0.558***		0.545***		0.566***	0.544***
	(0.0208)		(0.0201)		(0.0258)	(0.0295)
ln(Sales)		0.438***		0.429***		
		(0.0167)		(0.0162)		
Year FE	No	No	Yes	Yes	No	No
Industry FE	No	No	Yes	Yes	No	No
Industry-Year FE	Yes	Yes	No	No	Yes	Yes
Observations	35,992	35,992	36,768	36,768	35,992	35,992

Notes: Column (1) in this table contains the estimate for the coefficient β_S in the regression specified in equation (35). In Column (2), we replace the natural logarithm of employment with the natural logarithm of sales. In Columns (3) and (4), we replace industry-year fixed effects with industry and year fixed effects. In Column (5), the dependent variable is the number of unrelated startup acquisitions, and in Column (6), the number of related startup acquisitions. In all specifications, the dependent variable is normalized to have mean 1. * significant at 10%; ** significant at 5%; *** significant at 1%.

In Column (2), we use sales rather than employment as a measure of firm size, and continue to find a strong positive correlation. In Columns (3) to (4), we replace industry-year fixed effects with industry and year fixed effects. Finally, we consider as dependent variables the number of unrelated startup acquisitions (i.e., acquisitions of startups in a different 4-digit NAICS industry, Column (5)) and the number of related startup acquisitions (i.e., acquisitions of startups in the same 4-digit NAICS industry, Column (6)), finding that both are strongly positively correlated with acquirer firm size.

B.6 Drug development: data and additional results

B.6.1 Data construction

This section describes the construction of our dataset used for the empirical analysis in Section 3.4 of the main text. Our primary data source is Pharmaprojects, a database complied by the private company Citeline, which tracks the history of drug development projects.[48] We clean the raw data following the steps suggested by Cunningham *et al.* (2021), and then merge it with our SDC Platinum database on acquisition deals.

[48] See `https://www.citeline.com/en/products-services/clinical/pharmaprojects`.

Raw Data Import and Processing Pharmaprojects collects data at the drug level. For each drug contained in the dataset, we build a longitudinal history of events (containing e.g. the date on which the drug was added, the dates of clinical trials, drug launches, etc.). We also record each drug's therapeutic class (TC) and mechanism of action (MOA), converting the textual descriptions of these variables into a set of dummies.

Next, Pharmaprojects provides the name of the "originator" for each drug, i.e., the firm that started development. In 83% of all cases, this originator remains the only firm in control of the product. However, in other cases, drug ownership is transferred through mergers, acquisitions or transfers. To correctly identify the owner of a drug, we use the replication codes of Cunningham *et al.* (2021), which implement a text-mining algorithm on the Pharmaprojects variables "Overview" and "LatestChange". The authors use a regular expressions algorithm (`regexm`) to extract firm names associated with specific phrases indicating a transfer of ownership, such as "XYZ ACQUIRED BY ABC", "ACQUISITION OF XYZ BY ABC", "MERGER OF..." or "DEVELOPED BY XYZ (NOW ABC)". If the text-mining algorithm identifies a new owner, we assign the drug to this firm. Otherwise, we identify the "originator" as the owner.

Finally, to enable a merge with SDC Platinum, we manually clean all company names, removing legal entity identifiers (such as "INC", "LTD", "CORP", or "GMBH") or mapping subsidiaries and name variations to their ultimate parent or standardized name, and apply the `stnd_compname` name standardization routine developed by Wasi and Flaaen (2015).

Sample Selection To replicate the sample selection criteria from Cunningham *et al.* (2021), we impose the following restrictions:

- We restrict the sample to projects where the event "Drug Added" occurred between 1989 and 2010.
- We exclude "phantom" projects that show no development activity after initiation. That is, a project is included if it experiences at least one "Active Development" event. We define active development as the occurrence of any of the following: *New Launch, New Approval, Lead Identified, Development Restarted, First Launch, First Approval, New Mechanism, Nonproprietary Name Granted, Chemical Structure Reported, New Disease, New Licensee, New Patent, New Therapeutic Class, Novel Target Identified, Orphan Drug Status Granted* or *New Target*.

Applying these filters reduces the sample to 15,023 drug projects. These track the development of drugs from their first year in Pharmaprojects until 2017 (i.e. the last year in Cunningham *et al.*, 2021) or earlier if Pharmaprojects stopped tracking before. As an example, CitelineDrugID=8 is tracked in Pharmaprojects between 1992 and 2007. During this time it had 3 events of active development (1994, 1995, 1996). Therefore, our panel for this drug will have 16 observations for every year between 1992 and 2007, with no missing years in the middle.

Acquisition data and matching To track acquisitions, we continue to use SDC Platinum. However, we now restrict the sample to deals completed between 1981 and 2017, and deals in which the target firm belongs to the health care, medical devices, pharmaceutical, or chemical engineering industries. This is defined by a comprehensive list of SIC codes taken from Cunningham *et al.* (2021), including 2830-2836 (Drugs), 3840-3851 (Medical Instruments), and 8000-8099 (Health Services). This selection process yields 14,409 relevant M&A transactions where the target is a medical/pharmaceutical firm.

We then merge Pharmaprojects with SDC Platinum by using a multi-step fuzzy matching algorithm, supplemented by extensive manual verification. We first match the standardized Pharmaprojects company names to the SDC target names using the `reclink` command with a minimum matching score of 0.75. To eliminate false positives and increase the match rate, we performed manual review on all of candidate matches.[49]

B.6.2 Additional empirical results

In this section, we provide two additional results with respect to the main text. First, we replicate the main result in Cunningham *et al.* (2021), showing that there is a negative effect of acquisition on drug development if the target and acquirer drug portfolios overlap. To this end, we estimate the following regression:

$$\begin{aligned} \text{Development}_{it} = &\beta_1 \text{Acquired}_i + \beta_2(\text{Acquired}_i \times \text{Post}_{it}) + \beta_3(\text{Acquired}_i \times \text{Overlap}_i) \\ &+ \beta_4(\text{Acquired}_i \times \text{Post}_{it} \times \text{Overlap}_i) + \alpha_{FE} + u_{it}, \end{aligned} \quad \text{(B.1)}$$

where the dependent variable Development_{it} is a dummy variable taking value 1 when a drug i has a development event in year t. Acquired_i indicates whether a given drug i is acquired during the sample period, and Post_{it} indicates whether the drug-year observation (i, t) is posterior to the acquisition. The dummy variable Overlap_i switches to 1 if the acquiring firm has an existing product in the same therapeutic class (TC) that uses the same mechanism of action (MOA) as the acquired drug project. Standard errors are clustered at

[49] In order to replicate the results of Cunningham *et al.* (2021), we need to compute the overlap between acquirer and target drug projects. Thus, we perform a similar matching process between Pharmaprojects companies and SDC acquirer names.

the drug project level.

Table B.4 presents the results from this regression. Column (1), which does not include fixed effects, finds a negative and statistically significant estimate for the coefficient on the triple interaction term: acquired drug projects that overlap with the acquirers' portfolio are 4.39 percentage points less likely to have a development event in the post-acquisition years compared to non-overlapping acquired projects.

Column (2) includes project age and vintage fixed effects, in order to compare drugs that are created in the same year and are at the same stage of development. Column (3) adds drug project fixed effects to account for variation due to unobservable drug-specific characteristics. Finally, column (4) adds fixed effects for the interaction of age, therapeutic class and mechanism of action, controlling for heterogeneity in the life cycle of drugs targeting different diseases. Throughout these columns, we continue to find negative and significant estimates for the coefficient on the triple interaction.

Table B.4: Killer Acquisitions: replicating Cunningham *et al.* (2021)

	(1)	(2)	(3)	(4)
Acquired	-0.00427	0.0187***		
	(0.00572)	(0.00570)		
Acquired × Post	0.0148	-0.0257**	-0.0371**	-0.0240
	(0.00987)	(0.0100)	(0.0162)	(0.0218)
Acquired × Overlap	-0.0126	-0.00901		
	(0.0116)	(0.0114)		
Acquired × Post × Overlap	-0.0439*	-0.0469*	-0.0729*	-0.118**
	(0.0256)	(0.0257)	(0.0441)	(0.0590)
Age FE	No	Yes	Yes	Yes
Vintage FE	No	Yes	Yes	Yes
Project FE	No	No	Yes	Yes
Age × TC × MOA	No	No	No	Yes
Observations	113,159	113,159	113,061	83,889

Notes: We use a OLS estimator. The dependent variable is a dummy taking value 1 when a drug i has a development event in year t. "Acquired" indicates whether a given drug is acquired during the study period, and "Post" indicates whether the drug-year observation is posterior to the acquisition. "Overlap" is a dummy variable taking value 1 when the acquired drug overlaps with the product portfolio of the acquirer. Standard errors are clustered at the drug project level. * significant at 10%; ** significant at 5%; *** significant at 1%.

Finally, we re-estimate the same specifications as in Table B.4, limiting the treatment group to startups (i.e., targets below the median target age at the time of the acquisition), thus dropping drugs belonging to older acquired targets from the sample. Table B.5 presents the results. In this smaller sample, we have less statistical power, but the point estimates on the triple interaction are generally similar to the ones obtained in the full sample.

Table B.5: Killer acquisitions for startup targets

	(1)	(2)	(3)	(4)
Acquired	0.0177	0.0234		
	(0.0165)	(0.0169)		
Acquired x Post	-0.0161	-0.0296	-0.0256	-0.0483
	(0.0217)	(0.0218)	(0.0288)	(0.0403)
Acquired x Overlap	-0.0514*	-0.0503*		
	(0.0302)	(0.0301)		
Acquired x Post x Overlap	-0.0483	-0.0608	0.00158	-0.0684
	(0.0610)	(0.0609)	(0.0893)	(0.102)
Age FE	No	Yes	Yes	Yes
Vintage FE	No	Yes	Yes	Yes
Project FE	No	No	Yes	Yes
Age x TC x MOA	No	No	No	Yes
Observations	102,563	102,563	102,474	74,988

Notes: We use a OLS estimator. The dependent variable is a dummy taking value 1 when a drug i has a development event in year t. "Acquired" indicates whether a given drug belongs to a startup and is acquired during the study period, and "Post" indicates whether the drug-year observation is posterior to the acquisition. "Overlap" is a dummy variable taking value 1 when the acquired drug overlaps with the product portfolio of the acquirer. * significant at 10%; ** significant at 5%; *** significant at 1%.

C Calibration details and further robustness checks

C.1 Computing moments in the model

The **share of unrelated startup acquisitions** is equal to

$$\text{Moment}_1 = \frac{\gamma s_U}{\sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m(\omega,a)\Big(\gamma s_U + (1-\gamma)s_R(\omega,a)\Big)}. \tag{C.1}$$

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

$$\text{Moment}_2 = \frac{x_S \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m(\omega,a)\Big(\gamma s_U + (1-\gamma)s_R(\omega,a)\Big)}{1/\overline{n}}, \tag{C.2}$$

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 simulate a large cohort of firms, as described in Section C.2.

To match our regression evidence on the **effect of acquisitions for the implementation probability of startup ideas**, we proceed as follows. In our model, we know the implementation probability for each idea, depending on whether it is acquired or not. Thus, we can compute the causal effect of acquisition on 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

$$\beta_R \equiv \sum_{\omega\in\Omega}\sum_{a=1}^{+\infty} m_R(\omega,a)\Big[\ln\big(i_R(\omega,a)\big) - \ln\big(i_S(\omega)\big)\Big], \tag{C.3}$$

$$\beta_U \equiv \sum_{\omega\in\Omega} \phi(\omega)\Big[\ln\big(i_U(\omega)\big) - \ln\big(i_S(\omega)\big)\Big] = \frac{1}{\theta-1}\ln\left(\frac{\kappa_S}{\kappa_I}\right), \tag{C.4}$$

where $m_R(\omega,a) \equiv \frac{m(\omega,a)s_R(\omega,a)}{\sum_{\omega'}\sum_{a'} m(\omega',a')s_R(\omega',a')}$ denotes the distribution of acquired related startups over spending shares and technology gaps.[50] The overall causal effect is a weighted average

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

of these two statistics, with the weight given by the share of unrelated acquisitions (our first moment).

$$\text{Moment}_3 = \text{Moment}_1 \cdot \beta_U + (1 - \text{Moment}_1) \cdot \beta_R. \tag{C.5}$$

Our baseline target for this moment is zero. To make relative deviations from this moment well defined, we target $\exp(\text{Moment}_3)$ instead of Moment_3 in our calibration.

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

$$\text{Moment}_4 = \text{Moment}_1 \cdot \text{AcqPrem}_U + (1 - \text{Moment}_1) \cdot \text{AcqPrem}_R, \tag{C.6}$$

$$\text{with AcqPrem}_R \equiv (1 - \widetilde{\alpha}_R) \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m_R(\omega, a) \max \left[\frac{b_R(\omega, a)}{\kappa_S(\theta - 1)(i_S(\omega))^\theta}, 0 \right], \tag{C.7}$$

$$\text{AcqPrem}_U \equiv (1 - \widetilde{\alpha}_U) \sum_{\omega \in \Omega} \phi(\omega) \max \left[\frac{b_U(\omega)}{\kappa_S(\theta - 1)(i_S(\omega))^\theta}, 0 \right]. \tag{C.8}$$

For the **relationship between acquisitions and firm size**, we use our panel of simulated firms (described in Section C.2), to compute the yearly acquisition frequency for each simulated firm f and year t as

$$\text{AcqFreq}_{f,t} \equiv x_S \left(n_{f,t} \gamma s_U + (1 - \gamma) \sum_{j=1}^{n_{f,t}} s_R(\omega_{j,t}, a_{j,t}) \right) \tag{C.9}$$

where $n_{f,t}$ 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 frequency (the counterpart of regression (35) in the data) by running the OLS regression

$$\text{AcqFreq}_{f,t} = \alpha + \beta_S \ln \left(\text{Emp}_{f,t} \right) + u_{f,t}, \tag{C.10}$$

where $\text{Emp}_{f,t} = \sum_{i=1}^{n_{f,t}} \frac{\omega_i}{\lambda^{a_{it}}}$ is the firm's (normalized) employment, and $\text{Moment}_5 = \beta_S$.

The **probability of startup implementation conditional on no acquisition** is

$$\text{Moment}_6 = \frac{\sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) \Big(\gamma (1 - s_U) + (1 - \gamma) (1 - s_R(\omega, a)) \Big) i_S(\omega)}{1 - \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) \Big(\gamma s_U + (1 - \gamma) s_R(\omega, a) \Big)}. \tag{C.11}$$

This is the ratio of the arrival rate of startup ideas successfully implemented by startups, and the arrival rate of non-acquired startups.

The **aggregate growth rate** g is Moment$_7$. To decompose this growth rate, we divide aggregate innovation into three components:

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

where $g = \ln(\lambda)(I_{\text{Inc}}^{\text{CD}} + I_{\text{Inc}}^{\text{OwnProd}} + I_{\text{Ent}})$. The **share of growth accounted for by entrants** is

$$\text{Moment}_8 = \frac{I_{\text{Ent}}}{I_{\text{Inc}}^{\text{CD}} + I_{\text{Inc}}^{\text{OwnProd}} + I_{\text{Ent}}} \tag{C.12}$$

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

$$\text{Moment}_9 = \frac{I_{\text{Inc}}^{\text{CD}}}{I_{\text{Inc}}^{\text{CD}} + I_{\text{Inc}}^{\text{OwnProd}}}. \tag{C.13}$$

Finally, the BGP **exit rate of incumbents** is equal to the entry rate, given by

$$\text{Moment}_{10} = \frac{x_S}{1/\bar{n}} \left[\sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} m(\omega, a) i_S(\omega) \Big(\gamma (1 - s_U) + (1-\gamma)(1 - s_R(\omega, a)) \Big) \right]. \tag{C.14}$$

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

C.2 Simulating a cohort of firms

To compute some of the moments described above, 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 200 years. For each firm, we track its age and its product portfolio (i.e., the number of products and each product's spending share ω and technology gap a). 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.[51]

To carry out the simulation, we assume that all entrants start off with one product with technology gap $a = 1$. The share of entrants with a low-spending-share product is set to match the equivalent number on the BGP. Then, we update the product portfolio of each

[51] This simulation approach follows Garcia-Macia *et al.* (2019).

firm by using the product-level transition rates implied by the BGP policy functions.

C.3 Global identification

To analyze whether our targeted moments identify the internal parameters, we use a method first developed by Daruich (2026).

Denote our vector of internal parameters by $\boldsymbol{\theta} \in \mathbb{R}^{10}_{+}$. First, we create a 10-dimensional hyper-cube $\overline{\mathbb{P}}$ in the parameter space, choosing lower and upper bounds on each parameter. Then, we iteratively pick quasi-random realizations from $\overline{\mathbb{P}}$ using a Sobol sequence, which successively forms finer uniform partitions of the space. For a sufficiently large number of Sobol draws, this routine efficiently and comprehensively covers $\overline{\mathbb{P}}$. For each parameter draw, we solve the model and store its results in a matrix. After N Sobol draws (in practice, $N \approx 4.5$ million), we have a $N \times 10$ matrix $\boldsymbol{R}$ of results and a $N \times 10$ matrix $\boldsymbol{P} \in \overline{\mathbb{P}}$ 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.1 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). All in all, we find that all the parameters of the model are very well-identified by all four criteria.

Figure C.1: Global identification results.

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 lines marks the empirical target (red dotted) and the model's prediction (dashed blue). 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 $\widetilde{\kappa}_I = \kappa_I \cdot 0.1^{\theta}$.

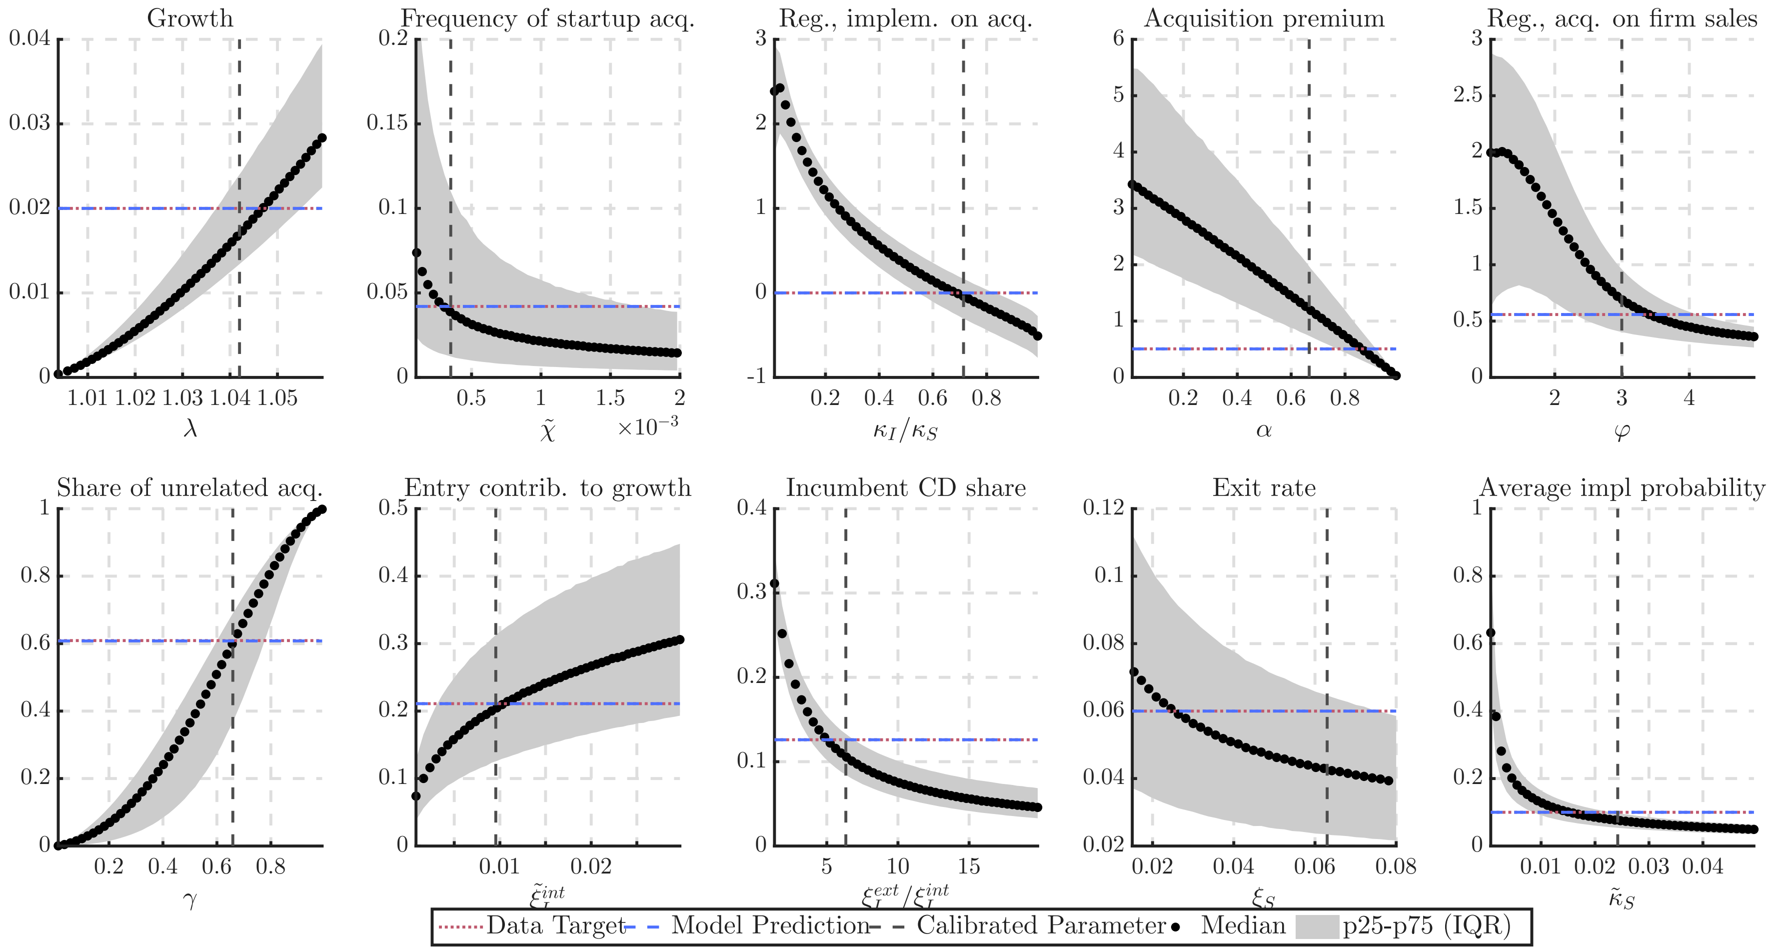

C.4 Quantitative results: Robustness exercises

Figures C.2 and C.3 explore the robustness of our main results with respect to all calibration targets that are not discussed in the main text. They indicate that overall, our main results are robust to reasonable variations in all moments.

Figure C.2: Robustness to other calibration targets (1/2)

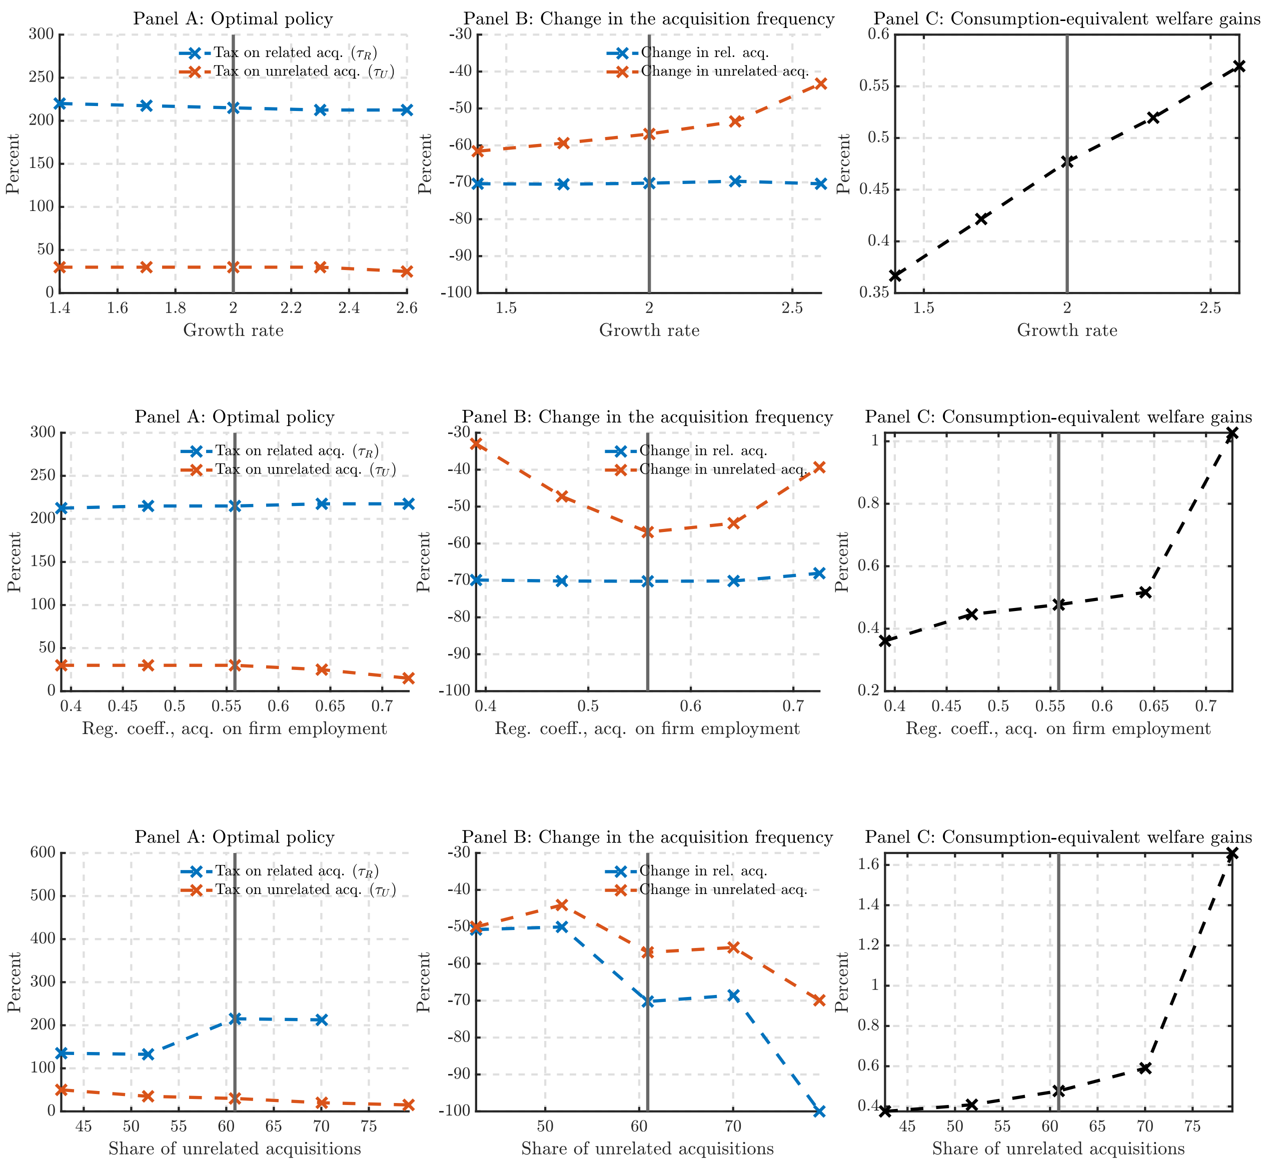

Notes: When targeting a share of unrelated acquisitions of 79%, the optimal tax on related acquisitions becomes infinitely high. This is the reason for which this tax rate is not shown in the left panel of the figure.

Figure C.3: Robustness to other calibration targets (2/2)

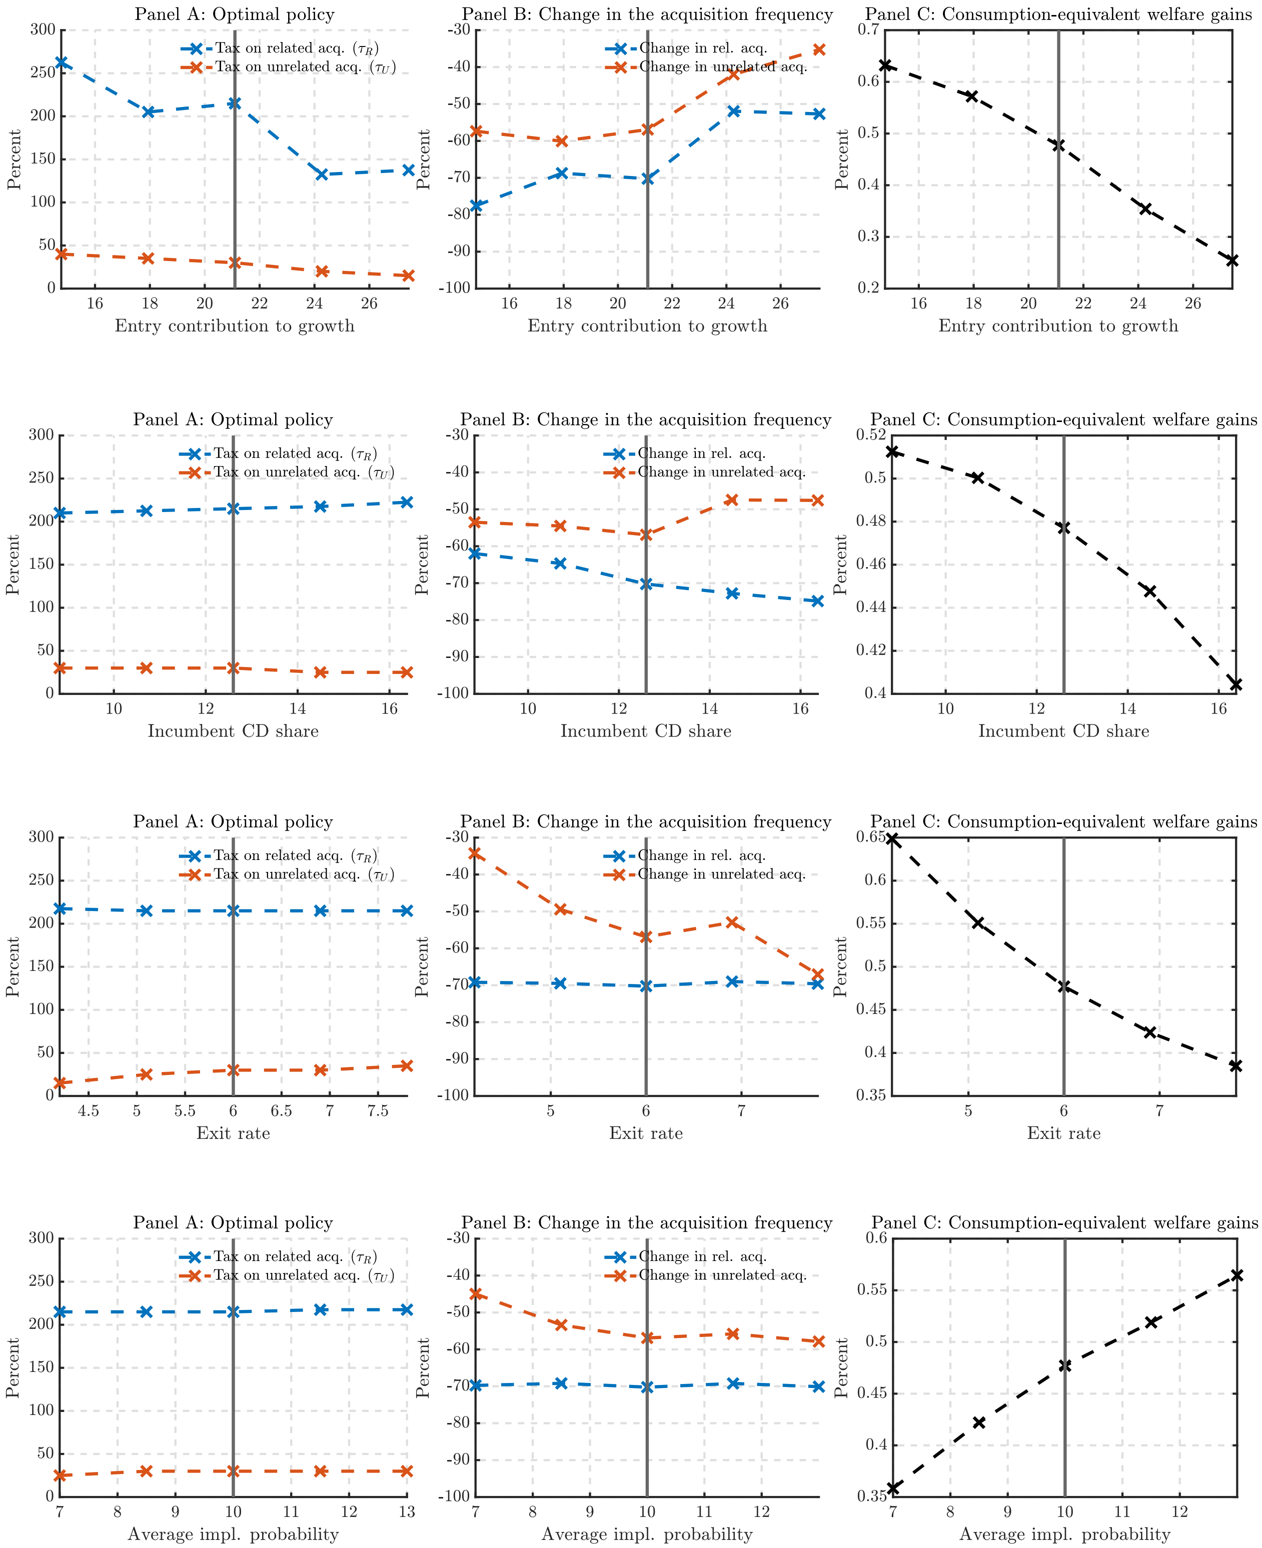

D Model extensions

In this section, we present and analyze two extensions of our baseline model.

D.1 Heterogeneous step sizes

Prior research has shown that startups have more radical ideas than incumbents, and that these ideas can lead to greater productivity increases (see e.g. Akcigit and Kerr, 2018). In our baseline model, all innovations increase productivity by the same amount. To match the observed heterogeneity between startup and incumbent ideas, this section develops an extended model in which startup ideas represent on average greater productivity increases.

Assumptions We continue to assume that an incumbent's idea represents a one-step advance on the productivity ladder. However, we now assume that a startup idea represents an advance of a_S steps, where the number of additional steps with respect to an incumbent idea $(a_S - 1)$ is stochastic, and drawn from a Poisson distribution with parameter ς.[52] This idea quality is revealed after all acquisitions and implementation decisions have been made.

Equilibrium conditions In this model, the value of a product in state (ω, a) is still given by equation (7) in the main text. However, the continuation values for unrelated acquisitions and for the appearance of a related startup (v_U and v_R) are now different.

Consider first the case of unrelated acquisitions. The outside option of a startup with an idea on a product with spending share ω' is now the value of the problem

$$\max_{i_S} \left\{ i_S \sum_{a_S=1}^{+\infty} \varrho(a_S) v(\omega', a_S) - \kappa_S i_S^{\theta} \right\}, \tag{D.1}$$

where $\varrho(a_S) \equiv e^{-\varsigma} \frac{\varsigma^{a_S-1}}{(a_S-1)!}$ is the probability that the startup's idea advances productivity by a_S steps. This implies that the startup chooses an implementation probability equal to

$$i_S(\omega') = \left(\frac{\sum_{a_S=1}^{+\infty} \varrho(a_S) v(\omega', a_S)}{\kappa_S \theta} \right)^{\frac{1}{\theta-1}}. \tag{D.2}$$

When an acquisition occurs, the continuation value of the incumbent, conditional on an acquisition offer $p_U^A(\omega')$, is the value of the problem

[52] This implies that, on average, a startup idea represents ς more steps on the productivity ladder than an incumbent idea. When $\varsigma = 0$, we recover the baseline model.

$$\max_{i_U} \left\{ i_U \sum_{a_S=1}^{+\infty} \varrho(a_S) v(\omega', a_S) - \kappa_I i_U^{\theta} \right\} - (1+\tau_U)\, p_U^A(\omega') \tag{D.3}$$

so the incumbent implements the acquired idea with probability

$$i_U(\omega') = \left(\frac{\sum_{a_S=1}^{+\infty} \varrho(a_S) v(\omega', a_S)}{\kappa_I \theta} \right)^{\frac{1}{\theta-1}}. \tag{D.4}$$

As was the case for the startup, this implementation probability only depends on the product's spending share ω'. It is then easy to see that the solution to the Nash bargaining problem between the incumbent and the startup is still given by equation (16) in the main text. Equations (15), (17) and (18) continue to hold as well, except $i_S(\omega')$ is defined by (D.2) instead of (11), and $i_U(\omega')$ is defined by (D.4) instead of (13).

For related acquisitions, we follow the same logic to show that the incumbent implements the idea of a related startup with probability

$$i_R(\omega, a) = \left(\frac{\sum_{a_S=1}^{+\infty} \varrho(a_S) v(\omega, a + a_S) - v(\omega, a)}{\kappa_I \theta} \right)^{\frac{1}{\theta-1}}. \tag{D.5}$$

Equations (22), (23), (24), (25) and (26) in the main text are still valid, but $i_R(\omega, a)$ is now given by equation (D.5) instead of (20), and $i_S(\omega')$ is given by equation (D.2) instead of (11). Similarly, the free entry condition for startup creation is still given by equation (28).

The BGP growth rate in the extended model holds

$$g = \ln(\lambda) \left(x + \sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} \left[m(\omega, a) \omega \left(z(\omega, a) + (1+\varsigma) x_S i(\omega, a) \right) \right] \right) \tag{D.6}$$

where $i(\omega, a)$ is still defined by equation (31). The growth formula is virtually the same as in the baseline, except that the contribution of startup ideas is now shifted by a factor $(1+\varsigma)$: as startup ideas are on average more radical, they contribute more to growth.

Finally, the flow equations characterizing the distribution of technology gaps change as follows. For any technology gap $a > 1$ and spending share ω, we have

$$\begin{aligned} & z(\omega, a-1) m(\omega, a-1) + x_S \gamma s_U i_U(\omega) \varrho(a) \phi(\omega) \\ & \quad + x_S \sum_{k=1}^{+\infty} \left[\left(1 - \gamma s_U - (1-\gamma) s_R(\omega, k) \right) i_S(\omega) \varrho(a) \right] m(\omega, k) \end{aligned}$$

$$+ x_S \sum_{k=1}^{a-1} \left[(1-\gamma)\, s_R(\omega,k) i_R(\omega,k) \varrho(a-k) \right] m(\omega,k) \tag{D.7}$$

$$= \Big(z(\omega,a) + x_S i\,(\omega,a) + x \Big) m(\omega,a).$$

This equation equates inflows into a state (the left-hand side) to outflows from that state (the right-hand side). Inflows into state (ω, a) are given by incumbent innovations in state $(\omega, a-1)$, startup innovations of size a, and startup ideas of size $a = k$ implemented by incumbents in state (ω, k). Once again, the distribution holds $\sum_{a=1}^{+\infty} m(\omega,a) = \phi(\omega)$. Together with the set of equations defined by (D.7), this condition gives a linear system that we can solve to find the invariant distribution m.

Results To analyze this extended model, we calibrate it to the data. The extended model has one additional parameter, ς, which pins down the quality difference between incumbent and startup ideas. We calibrate this parameter externally, following Akcigit and Kerr (2018), who find that external innovations by startups have on average a 35% greater impact on productivity than internal innovations by incumbent firms. Accordingly, we set $\varsigma = 0.35$. We keep all other external parameters at their baseline values, and re-calibrate the internally set parameters to match the same targets as in the baseline calibration. Table D.1 summarizes the fit of the extended model, showing that it matches the data just as well as the baseline.

Table D.1: Parameter values and model fit: Model extension with better startup ideas

Parameter	Value	Interpretation	Source / Moment	Data	Model
ρ	0.03	Discount rate	Standard value		
ϕ_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	1.033	Innovation step size	Growth rate	2.00%	2.00%
χ	0.376	Search cost scale	Frequency of startup acq.	4.20%	4.19%
κ_I/κ_S	0.725	Inc. relative impl. cost	Reg., implem. on acq.	0.000	-0.002
α	0.632	Inc. bargaining weight	Acquisition premium	50.7%	50.7%
φ	3.108	Search cost curvature	Reg., acq. on firm sales	0.558	0.557
γ	0.648	Share of unrelated startups	Share of unrelated acq.	60.9%	61.2%
ξ_I^{int}	0.576	Inc. int. innovation cost	Entry contrib. to growth	21.1%	21.0%
$\xi_I^{\text{ext}}/\xi_I^{\text{int}}$	7.078	Rel. ext. innovation cost	Incumbent CD share	12.6%	12.6%
ξ_S	0.066	Startup cost	Exit rate	6.0%	6.0%
κ_S	2.398	Startup implementation cost	Average impl probability	10.0%	10.0%

Notes: This table lists the parameter values for the calibration of the extended model described in Online Appendix D.1. For each internally calibrated parameter, the table lists the moment identifying the parameter, as well as the data and model value of that moment.

With this, we can determine the welfare-maximizing level of startup acquisition taxes. We proceed as in Section 5.3 and assume that the government announces tax rates τ_R and τ_U at $t = 0$. We then compute the transition to the new BGP implied by these taxes, and evaluate consumption-equivalent welfare along the transition path.

Figure D.1 (the equivalent of Figure 6 in the main text) illustrates our results. These are very similar to the baseline. The optimal policy imposes a tax of 172.5% on related acquisitions, and a tax of 25% on unrelated acquisitions (rather than 215% and 30% in the baseline). This triggers an increase in consumption-equivalent welfare of 0.35%, similar to the 0.48% gain in the baseline.

Figure D.1: The optimal acquisition tax - Model extension with better startup ideas

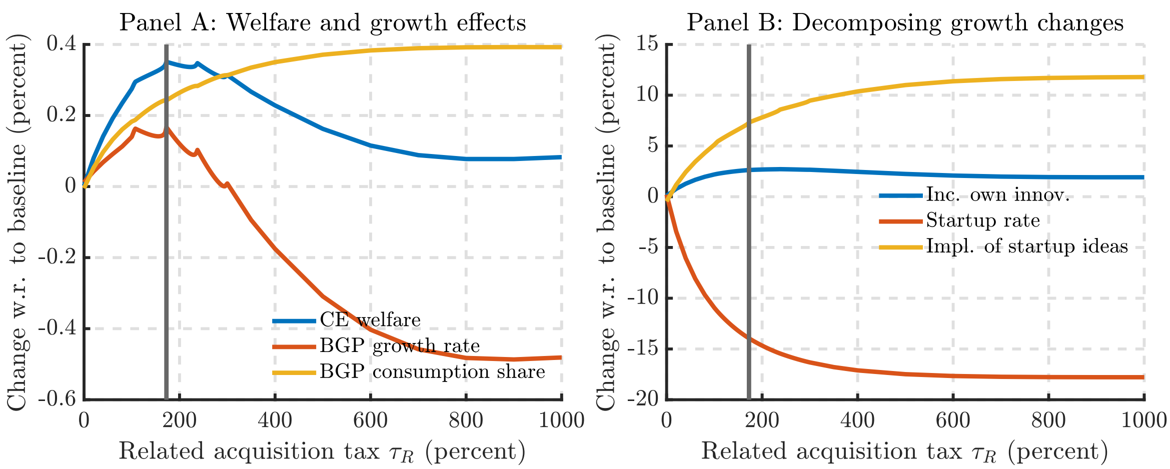

Intuitively, results are similar because both models target the same contribution of entry to overall productivity growth. To match this moment, the model with better startup ideas has slightly higher startup innovation costs: comparing Table D.1 and Table 5 in the main text, the extended model has both higher startup creation costs ξ_S and higher implementation costs for startup ideas κ_S. Thus, the economic return to startup ideas does not change much.

D.2 Acquisitions between incumbents

In our baseline model, incumbents could only acquire ideas from startups, not from other incumbents. While this is in line with our paper's focus on startup acquisitions, in this section, we extend the model to allow incumbents to acquire ideas from other incumbents.

Assumptions We now assume that the external innovations of incumbents are exactly analogous to startup ideas. Therefore, they now also need implementation investment, modeled exactly like the implementation investment for startups in the baseline. Moreover, with probability γ, they get matched to an unrelated incumbent (a randomly chosen firm from the continuum of incumbents), and with probability $1-\gamma$, to the related incumbent which currently produces the product to which the idea applies. Unrelated and related incumbents notice this idea with probabilities s_U and s_R (that is, we assume that the same search investment allows incumbents to notice both incumbents and startups).

Finally, we assume that acquisitions of startups and incumbent ideas can be taxed at different rates, to allow for differences in antitrust policy between these two cases.

Equilibrium conditions From the outset, we can note that an incumbent will never acquire an unrelated idea from another incumbent. Indeed, as shown in the main text, the only reason for unrelated acquisitions in our model are differences in implementation costs. As incumbents all share the same implementation cost κ_I, there is no surplus in transferring an unrelated idea from one incumbent to another. Thus, we can focus entirely on the acquisition of related ideas by incumbents.

In this new setup, we can rewrite the product-level HJB equation of an incumbent as

$$\rho v(\omega,a) = \max_{\substack{z,x\\ s_R,s_U}} \Bigg\{ \omega(1-\lambda^{-a}) \; - \xi_I^{\text{int}} z^{\psi} - \xi_I^{\text{ext}} x^{\psi} - \chi\left(s_R^{\varphi} + s_U^{\varphi}\right) + z\Big[v(\omega,a+1) - v(\omega,a)\Big]$$

$$+ x \sum_{\omega'\in\Omega} \sum_{a'=1}^{+\infty} m(\omega',a') v_X(\omega',a') + \tilde{x} v_R^I(s_R,\omega,a) + x_S v_R^S(s_R,\omega,a) + x_S \gamma v_U(s_U) \Bigg\}. \tag{D.8}$$

In this equation, $v_X(\omega',a')$ stands for the continuation value of an incumbent idea on a product with characteristics (ω',a'). $v_R^I(s_R,\omega,a)$ is the continuation value for an incumbent when another incumbent has an idea on its product, and $v_R^S(s_R,\omega,a)$ (denoted by $v_R(s_R,\omega,a)$ in the baseline model) is the continuation value for an incumbent when a startup has an idea on its product.

The internal innovation rate z, the change in value when an incumbent gets matched to an unrelated startup v_U, the unrelated search effort s_U, and the change in value when a startup has an idea on an incumbent's product, v_R^S, are still given by the same expressions as in the baseline model, i.e., equations (8), (17), (18) and (24).

To derive the new elements, note that if an incumbent's external idea on a product with characteristics (ω',a') is not acquired, it implements it with probability

$$i_X(\omega') = \left(\frac{v(\omega',1)}{\kappa_I \theta}\right)^{\frac{1}{\theta-1}}. \tag{D.9}$$

Note that this is the same probability with which the incumbent would implement an acquired unrelated startup idea. Thus, the target incumbent's outside option in an acquisition negotiation is $\kappa_I\,(\theta-1)\,(i_X(\omega'))^{\theta}$.

Next, we turn to the potential acquirer. When the related idea is not matched to it, or when it does not notice the potential target incumbent, it faces displacement from the latter firm, with probability $i_X(\omega)$. Thus, its value decreases in expectation by $i_X(\omega)v(\omega,a)$. In case the potential acquirer incumbent is matched to the potential target incumbent, an acquisition can take place. The situation is then exactly analogous to the one in which the related incumbent notices a startup. Thus, an acquisition takes place if and only if the surplus $b_R^I(\omega)$ is non-negative, where

$$b_R^I(\omega,a) \equiv (1+\alpha\tau_R^I)\left[(\theta-1)\,\kappa_I\left(\frac{(i_R(\omega))^{\theta}}{1+\tau_R^I}-(i_X(\omega))^{\theta}\right)+\frac{i_X(\omega)v(\omega,a)}{1+\tau_R^I}\right]. \tag{D.10}$$

Note that the first summand is now unambiguously negative, as the idea itself is always worth more to the target than to the acquirer. Thus, acquisitions only take place to preserve the acquirer's rents. τ_R^I stands for the tax rate on acquisitions of ideas of related incumbents.

When the surplus is non-negative, the acquisition price is

$$\kappa_I\,(\theta-1)\,(i_X(\omega))^{\theta}+(1-\widetilde{\alpha}_R^I)b_R^I(\omega,a), \qquad \text{where } \widetilde{\alpha}_R^I \equiv \frac{\alpha(1+\tau_R^I)}{1+\alpha\tau_R^I}. \tag{D.11}$$

With this, we can write down the continuation values for both the potential acquirer and the potential target. For the former,

$$v_R^I(s_R,\omega,a) = -i_X(\omega)v(\omega,a)+s_R\,(1-\gamma)\,\widetilde{\alpha}_R^I\max\Big(0,b_R^I(\omega,a)\Big). \tag{D.12}$$

For the latter,

$$v_X(\omega',a') = \kappa_I\,(\theta-1)\,\big(i_X(\omega')\big)^{\theta}+s_R(\omega',a')\,(1-\gamma)\,\Big(1-\widetilde{\alpha}_R^I\Big)\max\Big(0,b_R^I(\omega',a')\Big). \tag{D.13}$$

Accordingly, the search probability for related ideas and the external innovation rate of incumbents hold

$$s_R(\omega,a) = \left((1-\gamma)\,\frac{x\widetilde{\alpha}_R^I\max\big(0,b_R^I(\omega,a)\big)+x_S\widetilde{\alpha}_R^S\max\big(0,b_R^S(\omega,a)\big)}{\chi\varphi}\right)^{\frac{1}{\varphi-1}}, \tag{D.14}$$

where $\widetilde{\alpha}_R^S$ corresponds to $\widetilde{\alpha}_R$ in the main text. Now, search for related ideas is motivated by threats from both incumbents and startups. In turn,

$$x = \left(\frac{1}{\xi_I^{\text{ext}} \psi} \sum_{\omega' \in \Omega} \sum_{a'=1}^{+\infty} m(\omega', a') v_X(\omega', a') \right)^{\frac{1}{\psi - 1}}. \tag{D.15}$$

Replacing all our results into equation (D.8), product value holds

$$v(\omega, a) = \frac{\pi(\omega, a) + (\psi - 1) \left(\xi_I^{\text{int}} z(\omega, a)^{\psi} + \xi_I^{\text{ext}} x^{\psi} \right) + \chi (\varphi - 1) \left(s_R(\omega, a)^{\varphi} + s_U^{\varphi} \right)}{\rho + x i_X(\omega) + x_S \left(\gamma s_U i_U(\omega) + (1 - \gamma s_U) i_S(\omega) \right)}. \tag{D.16}$$

The free entry condition for startups, in turn, is still given by equation (28). The two last elements that change with respect to the baseline are the growth rate and the law of motion for the joint distribution of spending shares and technology gaps. For the growth rate, we have

$$g = \ln(\lambda) \left(\sum_{\omega \in \Omega} \sum_{a=1}^{+\infty} \left[m(\omega, a) \omega \left(z(\omega, a) + x \Big((1 - \gamma)\, s_R(\omega, a) i_R(\omega, a) + \right.\right.\right.$$
$$\left.\left.\left. (1 - (1 - \gamma) s_R(\omega, a))\, i_X(\omega) \Big) + x_S i(\omega, a) \right) \right] \right), \tag{D.17}$$

where the only change with respect to the baseline is that external innovations now have to be implemented, either by the incumbent that created them or by another incumbent that bought them up.

For the technology gap distribution, we now have for every technology gap $a > 1$

$$\Big(z(\omega, a - 1) + (x_S + x)\, (1 - \gamma)\, s_R(\omega, a - 1) i_R(\omega, a - 1) \Big) m(\omega, a - 1) =$$
$$\Big(z(\omega, a) + x_S i\,(\omega, a) + x \Big((1 - \gamma)\, s_R(\omega, a) i_R(\omega, a) + (1 - (1 - \gamma) s_R(\omega, a))\, i_X(\omega) \Big) \Big) m(\omega, a).$$

The principle here is the same as in the baseline. Inflows are given by incumbent innovations on their own product, either generated by themselves or acquired from a startup or incumbent. Any innovation triggers an outflow. As before, the distribution must hold

$$\sum_{a=1}^{+\infty} m(\omega, a) = \phi(\omega), \tag{D.18}$$

which closes the model.

Results The extended model has the exact same set of parameters as the baseline. Thus, we keep all externally set parameters at their baseline values, and re-calibrate the internally set parameters to match the same targets as in the baseline calibration. Table D.2 summarizes the fit of the extended model, showing that it matches the data moments well.

Table D.2: Parameter values and model fit: Model extension with incumbent acquisitions

Parameter	Value	Interpretation	Source / Moment	Data	Model
ρ	0.03	Discount rate	Standard value		
ϕ_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	1.038	Innovation step size	Growth rate	2.00%	2.00%
χ	0.213	Search cost scale	Frequency of startup acq.	4.20%	4.22%
κ_I/κ_S	0.721	Inc. relative impl. cost	Reg., implem. on acq.	0.000	0.009
α	0.722	Inc. bargaining weight	Acquisition premium	50.7%	50.9%
φ	2.695	Search cost curvature	Reg., acq. on firm sales	0.558	0.561
γ	0.73	Share of unrelated startups	Share of unrelated acq.	60.9%	60.8%
ξ_I^{int}	0.733	Inc. int. innovation cost	Entry contrib. to growth	21.1%	21.1%
$\xi_I^{\text{ext}}/\xi_I^{\text{int}}$	0.278	Rel. ext. innovation cost	Incumbent CD share	12.6%	12.6%
ξ_S	0.043	Startup cost	Exit rate	6.0%	6.0%
κ_S	2.245	Startup implementation cost	Average impl probability	10.0%	9.9%

Notes: This table lists the parameter values for the calibration of the extended model described in Online Appendix D.2. For each internally calibrated parameter, the table lists the moment identifying the parameter, as well as the data and model value of that moment.

To analyze the properties of this extended model, we first consider the effects of taxes on startup acquisitions, while leaving the acquisitions of ideas from other incumbents untaxed. That is, we assume that the government announces at $t = 0$ positive taxes on related and unrelated startup acquisitions, τ_R^S and τ_U, but keeps taxes on the acquisition of related incumbent ideas τ_R^I at zero. Figure D.2 illustrates our results. Again, the optimal policy is a high tax on related startup acquisitions (equal to 140%), combined with a lower tax on unrelated startup acquisitions (equal to 30%). As in the baseline, this policy essentially maximizes the positive response of incumbent own innovation. The resulting increase in consumption-equivalent welfare is equal to 0.42%, close to the baseline 0.48% increase.

While these results show that the interaction effects between acquisitions of ideas by startups and incumbents are limited, the extended model opens up the question whether acquisitions of ideas from other incumbents should also be taxed. Thus, we conduct an additional exercise in which we assume that the government announces the same (positive) tax rate for all acquisitions of related ideas, $\tau_R^S = \tau_R^I$. We find that the optimal tax rates are virtually unchanged, with a 135% tax rate for related startup acquisitions and acquisitions of related incumbent ideas, and a 30% tax rate for unrelated acquisitions. However, this broader policy leads to somewhat higher consumption-equivalent welfare gains of 0.46%.

Figure D.2: The optimal acquisition tax - Model extension with incumbent acquisitions