Interest Rates, Innovation, and Creative Destruction

Jonathan Goldberg[*] David López-Salido Craig A. Chikis

Federal Reserve Board

November 2, 2020

Abstract

Do very low interest rates harm innovation, long-run growth, and competition? Using a quantitative endogenous growth model, we find that a lower interest rate boosts growth with the markup distribution little changed. Our framework nests canonical assumptions about creative destruction and the technology of R&D. To match the cross section of markups, profit volatility, R&D, and innovation output, as well as entrants' employment share and contribution to aggregate innovation, the model requires some "advantage of backwardness," with market laggards or entrants always having at least some chance of more-than-incremental innovation. A lower interest rate spurs innovation by raising the value of the profits from innovation; in general equilibrium, strategic interactions among firms and a rise in the wage dampen, but do not overturn, the valuation-driven increase in R&D and growth. To be sure, imposing severe (and counterfactual) restrictions on creative destruction leads to a growth "speed limit" at low interest rates, with growth declining as the interest rate falls. However, this decline in growth reverses under modest departures from these severe restrictions, if labor supply is elastic, if patent policy is adjusted optimally, or if credit access is restricted. The growth "speed limit" economy offers policymakers a free lunch: weakening patent protection boosts growth and reduces markup-related production distortions. Our results suggest that very low interest rates do not harm growth and competition unless the forces of creative destruction are unrealistically shackled.

Keywords: real interest rate, creative destruction, growth, markups, innovation.
JEL Codes: L16, O31, O34.

[*]Jonathan Goldberg (corresponding author, jonathan.goldberg@frb.gov). David López-Salido: david.lopez-salido@frb.gov. Craig A. Chikis: craig.a.chikis@frb.gov. We thank Gadi Barlevy, Ernest Liu, and Jane Olmstead-Rumsey for comments. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of anyone else associated with the Federal Reserve System.

1. Introduction

How do interest rates affect innovation, long-run growth, and competition? Very low interest rates encourage innovation by increasing the value of the profit stream from innovation, but can very low interest rates stifle growth and competition by entrenching market leaders? How should innovation policies—such as patent policy and R&D subsidies—respond to persistent changes in the interest rate, and how would such policy adjustments affect the relation of growth, interest rates, and markups? To address these questions, we develop a Schumpeterian multisector model à la Aghion-Howitt in which firms strategically engage in R&D to climb a productivity ladder.

For a firm considered in isolation (that is, holding constant all factors external to the firm, such as the wage and competitors' strategies), a decline in the discount rate used to make R&D decisions is expected to increase R&D and hence contribute to growth. However, strategic interactions among firms and general-equilibrium changes in the wage and competitiveness of the economy can, in principle, amplify, dampen, or even overturn this growth effect.

We investigate how these strategic and general-equilibrium forces are affected by creative destruction and the technology of R&D (how innovations are produced). Our model nests, at the extremes, innovation by market laggards and entrants that always builds incrementally on laggards' existing technologies ("slow catch-up") and always achieves parity with or even leapfrogs the incumbent leader (termed "quick catch-up" by Acemoglu and Akcigit (2012)). Entrant R&D can similarly result in slow or quick advances, be more or less productive than incumbent R&D, and be undirected (Klette and Kortum (2004)) or focused on a single industry (Acemoglu and Cao (2015)). As with innovation, a patent expiry can advance a laggard's technology incrementally, all the way to parity with the market leader, or an intermediate number of steps. Overall, these assumptions capture many dimensions of the "advantage of backwardness" (Gerschenkron (1962)).

These assumptions affect how market power varies across industries, how much R&D is conducted and by which types of firms, the volatility of profits for firms in different market positions, and the roles of entrants in contributing to growth and employment. Accordingly, to identify the model's parameters, we compute 14 key moments characterizing the cross section of markups (Hall (2018)), profit volatility, firm innovation output (Kogan et al. (2017)), and R&D. The model matches closely these key moments and performs well on young firms' employment share and the roles of incumbent innovation, reallocation, and entry in Foster et al. (2001)'s growth decomposition. The calibrated parameters depict an economy with a high degree of quick catch-up from entrant innovation and patent expiry, but slow catch-up through laggard innovation.

We study a decline in the household discount rate in this calibrated economy.[1] In the model that best fits the data, a lower discount rate leads to an increase in the growth rate and a decrease in the interest rate. The distribution of markups is little changed. To gauge whether entry is essential to our results, we then turn entry off and recalibrate the model. Without entrants, the model requires fairly quick catch-up through laggard innovation to match the data; the quality of fit is fairly good for non-entry moments. A lower interest rate increases growth and reduces the average markup.

Next, we restrict creative destruction, beyond turning off entry, by assuming extremely slow catch-up for laggards—that is, laggard innovation can only be incremental and patent expiry closes the technology gap between leader and laggard only slowly. To match the aggregate growth rate and average markup, we set the patent expiry rate to a very high level. As the discount rate falls, growth rises monotonically (as in our benchmark model) and the average markup rises. The monotonicity of the growth-interest rate relation holds for perfectly elastic and inelastic labor supply, and when competitors in each industry produce perfect and imperfect substitutes.

However, when we further restrict creative destruction by considerably reducing the patent expiry rate (i.e., increasing patent protection), we obtain a growth "speed limit"—there is a maximum growth rate such that further declines in the interest rate reduce growth. This result implies an inverted-U relation between growth and the interest rate.[2]

[1]Changes in the discount rate represents factors such as demographics and risk appetite that are not modeled and affect the demand for claims on firm profits, e.g. Eggertsson et al. (2018) and Liu et al. (2020).

[2]This result is consistent with the relation between growth and low interest rates in Liu et al. (2020). Our assumptions about creative destruction are different than theirs—even in our speed limit economy, we do not have "pure" slow catch-up through R&D *and* patent expiry—and our inverted-U arises through a different mechanism. These differences described in detail in Section 5.4.

In this economy, as the interest rate falls, prospective superstars' R&D rises sharply, as these firms seek to establish an insurmountable lead and discourage competitors. That is, a low interest rate supercharges the escape competition effect of Aghion and Howitt (1992). Starting from a high interest rate, declines in the interest rate therefore lead to a build-up of market power. However, prospective superstars' R&D increases the aggregate demand for labor and, if aggregate labor supply is inelastic, pushes up the wage and crowds out R&D by firms that already achieved superstar status. Because superstars have become predominant, this crowding out reduces aggregate growth. Thus, a (general-equilibrium) labor market effect underpins our "speed limit" result. Indeed, if we allow for perfectly elastic labor supply, the growth speed limit disappears.

The growth speed limit result is fragile in other ways as well. With only moderate departures from the limiting case of "pure" slow catch-up, the inverted-U relation vanishes. The inverted-U relation also depends on a low patent expiry rate that, combined with slow catch-up, implies a counterfactually sclerotic economy: profit volatility is extremely low, market power is very high, and the distribution of innovation output (Kogan et al. (2017)) completely lacks the fat right tail in the data. Moreover, were a social planner to set the patent expiry rate optimally (decreasing patent protection), markups would be much lower and the inverted-U relation would disappear.

To unpack these results, we develop a new methodology called the innovation multiplier. The innovation multiplier is a mapping from the cross section of individual firm R&D effects of a lower discount rate (holding constant all factors external to the firm) into the total effect on growth. The multiplier identifies—and quantifies—strategic interactions that reduce growth, and other strategic interactions that increase it. A decline in the interest rate boosts innovation by entrants, which dampens the rise in innovation by market leaders—innovation for market leaders is more profitable when the leader is less likely to be leapfrogged by an entrant. We call this channel the strategic trickle-down effect (Acemoglu and Akcigit (2012)). In contrast, in highly competitive industries, higher laggard innovation powerfully amplifies the rise in innovation by market leaders due to strategic escape competition (Aghion et al. (2001)). Growth depends not only on firm innovation strategies, but also the economy's competitiveness, because firms in more competitive industries innovate more. The multiplier embeds how the resulting composition (or extensive margin) effects and general-equilibrium adjustments in the wage impact growth.

Overall, we find that strategic interactions and wage and composition effects in our Schumpeterian model significantly dampen—but do not reverse—the rise in growth from a lower interest rate. If firm-level R&D were to adjust based only on the higher valuation of profits from innovation, the total effect on growth would be an 18 basis point increase in the long-run growth rate. However, factoring in strategic interactions, the change in the wage, and the composition effect, growth rises only 7 basis points. Growth-reducing strategic trickle-down effects of higher entry more than outweigh the growth-increasing strategic escape competition effects of higher R&D by laggards in competitive industries. The total effect on growth of a lower discount rate is slightly further reduced by a growth-dampening composition effect, arising from a slightly lower share of industries that are highly competitive, and a rise in the wage. In our benchmark economy, these growth-reducing forces on balance dampen growth, but are insufficient to lead to lower growth. In contrast, in our speed limit economy, the Jorgensonian increase in growth is completely reversed by these forces: a lower discount rate leads to lower growth, at least when starting from a low interest rate.

A lower interest rate could in principle give rise to a much higher average markup, because the rise in leader innovation is mostly unmatched by laggards. Indeed, the share of highly uncompetitive industries increases slightly. However, the average markup, while non-monotonic in the interest rate, is quantitatively little affected by the interest rate. Key to this result is a large increase in innovation by laggards in competitive industries and by entrants, which implies that lower interest rates increase the share of highly competitive industries. Thus, there is some "hollowing out" of the distribution of industries, with fewer industries with an intermediate level of competitiveness.

Next, we investigate how credit access can affect the relation of growth, markups, and the interest rate. Following Aghion et al. (2019a), we introduce a limited-commitment constraint affecting R&D. We find that the constraint binds for tied firms and for leaders in competitive industries. In the benchmark economy, reduced credit access reduces growth and markups, due to reduced escape-competition innovation by these firms in competitive industries. However, in the speed limit economy, at low interest rates, reduced credit access *increases* growth, by eliminating the growth-stifling composition effects that give rise to the growth speed-limit. Indeed, with credit access severely reduced, the inverted-U relation between growth and the interest rate vanishes.

We next study the welfare implications of innovation policy. We have five main results. First, in our benchmark model, the social planner prefers a lower patent expiry rate than the calibrated expiry rate that matches market-power moments. The optimal patent expiry rate boosts annual growth by 24 basis points, with somewhat worse markup-related production distortions. Second, as the interest rate falls, the social planner prefers to strengthen patent protection. Third, in the growth speed-limit economy, the opposite conclusions hold: the social planner prefers a higher patent expiry rate than the expiry rate that gives rise to the speed limit, and prefers to raise the expiry rate as the interest rate falls.

Strikingly, in the speed limit economy, there is a "free lunch" for the social planner: by weakening patent protection, the planner boosts growth by 17 basis points and reduces markup distortions, with the ratio of output-to-potential (the welfare-relevant measure of the absence of markup distortions) rising 2 percentage points. Moreover, by optimally adjusting the patent expiry rate in a low interest rate environment, the planner avoids the build-up of market power essential to the growth speed limit—and the growth speed limit disappears. Fourth, in the benchmark and the speed-limit economies, the social planner can boost welfare considerably (relative to the welfare that obtains with optimal uniform patent policy) by adopting a patent policy that gives stronger patent protection in less competitive industries. Such state-dependent policies (Acemoglu and Akcigit (2012)) allow for much higher growth at low interest rates. Fifth, an R&D subsidy applying only to laggards boosts growth by 10 basis points when set optimally and the welfare-maximizing subsidy rises as the interest rate falls.

Related Literature. Our paper contributes to the debate regarding recent trends of increasing markups, low productivity growth, and the rise of superstar firms.[3] Eggertsson et al. (2018) parse the roles of exogenous shocks to markups, productivity growth, and

[3]See Hall (2018), De Loecker et al. (2020), Fernald et al. (2017), and Autor et al. (2020), among many others.

discount rates. Aghion et al. (2019b) point to falling costs of spanning multiple markets as an underlying driver. Akcigit and Ates (2019)) emphasize declining knowledge diffusion. Our paper also contributes to the literatures on misallocation of inputs for production (Peters (2020)) and R&D (Acemoglu et al. (2018)). Our focus, distinct from these papers, is the role of interest rates and policy design in a low rate environment.

In an important contribution, Liu et al. (2020) study an economy with severe barriers to creative destruction. Specifically, there is no entry, and innovation by market laggards can advance the laggard's technology only very incrementally. Patent expiry also is assumed to lead to "pure" slow catch-up. In this economy, Liu et al. (2020) demonstrate theoretically a remarkable result: a very low interest rate leads to a low-growth economy with permanently entrenched leaders and effectively infinite markups. Liu et al. (2020) also calibrate their model to match the aggregate growth rate and the average markup and obtain a less extreme (but more realistic) version of their theoretical result. Our results are complementary to Liu et al. (2020): Rather than assume pure slow catch-up, we flexibly parametrize the forces of creative destruction and calibrate the model to match a rich array of cross sectional moments.[4]

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 provides a qualitative analysis. Section 4 describes our indirect inference approach and the parameter estimates. Section 5 presents the main quantitative results, which 6 unpacks using the innovation multiplier. Section 8 studies policy and welfare. Section 9 concludes.

2. Model

2.1 Preferences and final goods

The economy admits a representative household with a logarithmic utility function

$$\int_{t=0}^{\infty} e^{-\rho t} \ln(C(t))dt. \tag{1}$$

[4] Section 5.4 provides a detailed comparison with Liu et al. (2020).

$C(t)$ is consumption of the final good at time t and $\rho > 0$ is the discount rate. One unit of labor is supplied inelastically.[5]

The final good is produced using a continuum of intermediate goods according to the Cobb-Douglas production function, $\ln Y(t) = \int_0^1 \ln y_j(t)dj$, where $y_j(t)$ is the quantity of intermediate good j. The final good is the numeraire and is sold in a perfectly competitive market.

2.2 Intermediate goods market

Each industry j includes two firms. A firm i in industry j can produce the intermediate good j using a linear production technology, $y_{i,j} = q_{i,j}l_{i,j}$, where $q_{i,j}$ is the firm's labor productivity and $l_{i,j}$ is the amount of labor hired. The price of intermediate good j is p_j. Two small comments are in order. First, we drop time subscripts to describe the static allocation, or the price and quantity of each intermediate good taking each firm's productivity as given. Second, with a slight abuse of notation, the firm with the highest productivity is referred to using the subscript i and the other firm is referred to using $-i$.

We assume Bertrand competition, in which firms set prices. Each intermediate good is therefore sold at its limit price, $p_j = w/q_{-i,j}$, where w is the wage rate. Only the leader firm produces.[6] Because of the Cobb-Douglas technology for final good production, sales are equalized across industries, with $p_j y_j = Y = C$.

We use the Lerner index, or the ratio of price minus marginal cost to price, as a measure of market power (Hall (2018)). The Lerner index in industry j is

$$\mathscr{L}_j \equiv (p_j - \frac{w}{q_{i,j}})\frac{y_j}{p_j y_j} = 1 - \frac{q_{-i,j}}{q_{i,j}} = 1 - \psi_j^{-1}, \tag{2}$$

[5]These household preferences imply an elasticity of intertemporal substitution (EIS) equal to 1. Section 5.1 relaxes this assumption. Sections 5.2 and 5.3 study versions of this economy with elastic labor supply.

[6]These results follow from the assumption that the intermediate goods firms in each industry produce perfect substitutes. Section 5.2 recalibrates our economy under the assumption that the firms in each industry produce imperfect substitutes.

where ψ_j is the gross markup, or

$$\psi_j \equiv p_j \frac{q_{i,j}}{w} = \frac{q_{i,j}}{q_{-i,j}}. \tag{3}$$

The Lerner index and the markup in industry j are therefore functions of the productivity gap between the industry's leader and laggard. Labor demand is also a function of the industry markup, with $l_{i,j} = \frac{q_{-i,j}}{q_{i,j}} \frac{Y}{w} = (1 - \mathscr{L}_j)\frac{Y}{w} = \psi_j^{-1} \frac{Y}{w}$.[7] In a neck-and-neck industry, the two firms have the same productivity, $q_{i,j} = q_{-i,j}$, profits are zero, and the wage rate is equal to the marginal product of labor.

2.3 R&D, entry, and innovation

Following Aghion and Howitt (1992) and Grossman and Helpman (1991), firms in each industry are ordered on a quality ladder. Each rung represents a proportional productivity improvement of scale $\lambda > 1$. Leader productivity therefore exceeds laggard productivity by a factor $\lambda^{s(t)} = \frac{q_{i,j}(t)}{q_{-i,j}(t)}$, where s is the number of rungs separating the two firms at time t.[8] Correspondingly, the Lerner index is $\mathscr{L}_j(t) = 1 - \lambda^{-s(t)}$ and the gross markup is $\psi_j(t) = \lambda^{s(t)}$.

The *technology position* of a firm at time t is denoted by $s(t) \in \{-\bar{s}, \ldots, \bar{s}\} \equiv S$, where we assume that the maximum possible technology gap between leader and follower within an industry is $\bar{s}$.[9] For instance, a firm with position $s > 0$ is a leader s steps ahead of the follower. A firm with position $s < 0$ is a follower $|s|$ steps behind the leader. A firm with position $s = 0$ is tied. We also assume that the set of technology gaps is $s \in \{0, 1, ..., \bar{s}\} \equiv S^+$.[10]

Firm productivity $\{q_{i,j}(t), q_{-i,t}(t)\}$ and therefore the technology gap $s(t)$ can change for three reasons: innovation by an incumbent firm; patent expiry; and entry of a new firm.

[7] Industry profits and labor demand do not depend on the level of firms' productivity.

[8] The number of rungs between firms in industry j at time t is defined as $s_j(t)$. We generally omit the j and t indexes to simplify notation.

[9] With some abuse of notation, $s(t) \in \{0, ..., \bar{s}\}$ for an industry is the number of rungs between leader and laggard, while $s(t) \in \{-\bar{s}, ..., \bar{s}\}$ for a firm is its technology position. The assumption that the maximum possible industry gap is $\bar{s}$ implies that a leader cannot innovate further when it is $\bar{s}$ steps ahead. The purpose of this assumption is to make the state space finite. In our quantitative analysis, we set $\bar{s}$ to a high value. The properties of the balanced growth path of our calibrated model are very little changed if we double $\bar{s}$.

[10] That is, we assume that the initial condition of the economy features technology gaps $s \in \{0, 1, ..., \bar{s}\}$.

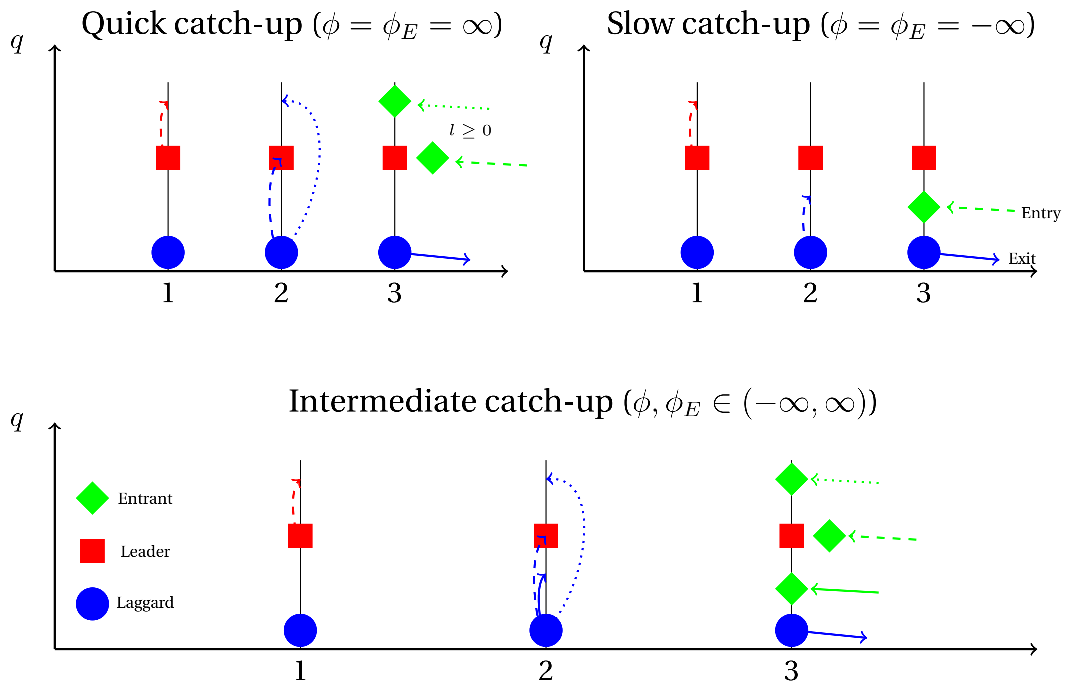

Figure 1: **Quality ladder.** Top left panel (QUICK CATCH UP): An innovating leader advances one rung (Line 1). An innovating laggard advances to a tied or leadership position depending on the leapfrogging parameter l (Line 2). An entrant arrives l_E steps ahead of the leader, displacing the laggard (Line 3). Top right panel (SLOW CATCH UP): Leader innovation is the same as with quick catch up (Line 1). An innovating laggard advances one rung (Line 2). An entrant arrives one step ahead of the exiting laggard (Line 3). Bottom panel: In the intermediate case between these extremes, an innovating laggard advances a stochastic number of rungs (Line 2). An entrant arrives a stochastic number of steps ahead of the displaced laggard (Line 3).

Figure 1 illustrates how firms move along the quality ladder.

Incumbent innovation. The innovation arrival rate of an incumbent firm in position $s \in \{-\bar{s}, ..., \bar{s}\} \equiv S$ is denoted $x_s(t)$. To achieve this innovation arrival rate, a firm hires $G(x_s(t); B) = [x_s(t)/B]^{\frac{1}{\gamma}}$ workers as R&D scientists. $B > 0$ is a scale parameter affecting the productivity of R&D scientists employed by incumbents. $\gamma > 0$ is a curvature parameter capturing the convexity of R&D costs in the arrival rate.[11] R&D expenditures are subsidized at rate $\tau_{R\&D}$.

[11]R&D costs $G(x; B)\omega$ are therefore increasing and convex in x and satisfy the Inada-type condition $G'(0) > 0$.

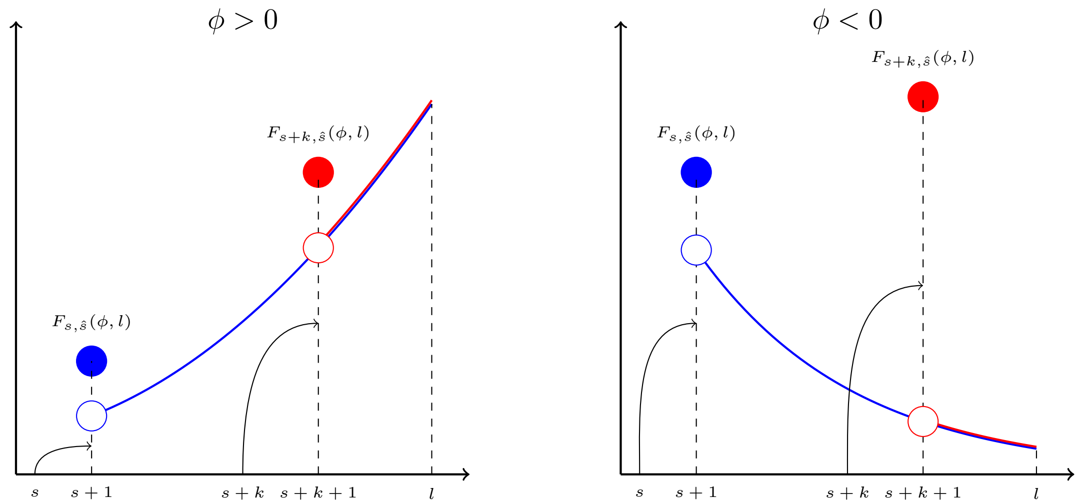

Figure 2: **Innovation advancement function.** An innovating incumbent in position $s < l$ arrives in position $\hat{s} > s$ with probability $F_{s,\hat{s}}$, shown in blue for $\phi > 0$ (left panel) and $\phi < 0$ (right panel). The distribution over the new position $\hat{s}$ for an innovating firm in position $s + k < l$ is shown in red. $l \in \{0, ..., \bar{s}\}$ is the leapfrogging parameter.

Upon innovating, a firm in position $s \in S$ moves one or more rungs up the quality ladder, obtaining the position $\hat{s}$ with probability $F_{s,\hat{s}}$. Thus, an innovating firm with productivity $q(t)$ obtains a new productivity $q(t + \Delta t) = \lambda^{\hat{s}-s} q(t)$.

We specify the advancement function $F_{s,\hat{s}}$ to nest a variety of assumptions in the literature about technological progress. A leapfrog parameter $l \in \{0, ..., \bar{s}\}$ is the maximum technology position that a laggard can potentially obtain upon innovation. If $l = 0$, at most, an innovating laggard catches up to, but does not surpass, the productivity of the leader. If $l > 0$, an innovating laggard might surpass the leader on the quality ladder. A larger value of l implies the possibility of more "radical" innovation by laggards (e.g., Acemoglu and Cao (2015)).

Together with the leapfrogging parameter l, a second parameter ϕ captures the *speed* at which an innovating laggard moves ahead. A higher value of ϕ implies a greater probability of an innovating laggard jumping all the way to position l. For $\phi \in (-\infty, \infty)$, the advancement function $F_{s,\hat{s}}$ is defined as follows. A firm in position $s < l$ (that is, a laggard firm or a leader firm less than l steps ahead) has a probability of advancing to position $\hat{s}$

$$F_{s,\hat{s}}(\phi, l) = \begin{cases} f_0 \sum_{\tilde{s}=-\bar{s}}^{s+1} \exp(\phi \tilde{s}) & \text{if } \hat{s} = s+1 \\ f_0 \exp(\phi \hat{s}) & \text{if } \hat{s} > s+1 \text{ and } \hat{s} \leq l, \end{cases} \tag{4}$$

and $F_{s,\hat{s}} = 0$ otherwise. The constant f_0 ensures that $\sum_{\hat{s} \in \{-\bar{s},\ldots,\bar{s}\}} F_{s,\hat{s}} = 1$, $\forall s$. We assume that leaders at or beyond position l advance (only) one step at a time. That is, for $s \geq l$, $F_{s,s+1} = 1$. This setup provides enough flexibility to capture a range of assumptions about leader innovation, including the common assumption that leader innovation is incremental ($l = 0$). The advancement function F is illustrated in Figure 2.

Our specification of the advancement function F nests several approaches in the literature. The "quick catch-up" setup of Aghion et al. (2001), with laggards catching up with a single innovation but unable to leapfrog, corresponds to the case of $\phi = \infty$ and $l = 0$.[12] We nest two additional quick catch-up settings: leapfrogging as in Acemoglu and Akcigit (2012) is obtained under $\phi = \infty$ and $l = 1$; and "radical" innovation as in Acemoglu and Cao (2015) corresponds to $\phi = \infty$ and $l > 1$. With $\phi = -\infty$, we nest the "slow catch-up" assumption of Acemoglu and Akcigit (2012) and Liu et al. (2020), in which an innovating firm advances only one step at a time. With $\phi \in (-\infty, \infty)$, the step size is stochastic as in Akcigit and Kerr (2018). We accommodate a range of intermediate catch-up speeds, with $\phi = 0$ corresponding to a uniform probability of landing on each step between $s+2$ and l.

Our specification also embeds the "advantage of backwardness" (Gerschenkron (1962)). That is, the probability that an innovation advances a firm more than one step is (weakly) increasing as a firm falls further behind. The advantage of backwardness is controlled by the parameters ϕ and l. For a laggard in position $s < 0$, it is more likely to move ahead by $\tilde{s}+1$ steps than to move ahead by $\tilde{s}$ steps, for $\tilde{s} \in \{2, ..., l-s\}$, if and only if $\phi > 0$. Akcigit et al. (2018) also embed an advantage of backwardness, while allowing for the possibility of laggards jumping all the way to the maximum technological gap. We approximate their setup with $\phi \in (0, \infty)$ and $l = \bar{s}$.

[12] Formally, we define $F_{s,\hat{s}}(-\infty, l)$, with $F_{s,\hat{s}}(-\infty, l) = 1$ if and only if $\hat{s} = s+1$. We similarly define $F_{s,\hat{s}}(\infty, l)$, with $F_{s,\hat{s}}(\infty, l) = 1$ if and only if $\hat{s} = s + l$. Note that $F_{s,\hat{s}}(\phi, l)$ converges uniformly to $F_{s,\hat{s}}(-\infty, l)$ as $\phi \to -\infty$ and similarly $F_{s,\hat{s}}(\phi, l)$ converges uniformly to $F_{s,\hat{s}}(\infty, l)$ as $\phi \to \infty$.

Patent expiry. Patent expiry allows a laggard to partly or completely catch up with its competitor. A patent expiry in an industry with gap $s > 0$ shrinks the gap to $\hat{s} \in \{0, .., s-1\}$. The new gap $\hat{s}$ has probability distribution $F^p_{s,\hat{s}}$. For a leader in position s, a patent expiry implies a new position $\hat{s}$ with probability distribution $F^p_{s,\hat{s}}$. For a follower in position $s < 0$, the distribution over its new position is symmetric: the follower moves to position $\hat{s}$ with probability $F^p_{-s,-\hat{s}}$. If a patent expiry occurs at time t, the productivity of a laggard with productivity q at time t becomes $q(t + \Delta t) = \lambda^{(s-\hat{s})} q(t)$.[13]

The parameter ζ characterizes the patent transition function, with $F^p_{s,\hat{s}}(\zeta) = F_{-s,-\hat{s}}(\zeta, 0)$ $\forall s > 0$. The case $\zeta = \infty$ corresponds to quick catch-up patent expiry: an industry's technology gap closes completely, with certainty, upon the expiry of the leader's patent. This assumption is the most common in the literature (e.g., Akcigit and Kerr (2018)). In contrast, $\zeta = -\infty$ corresponds to extreme slow catch-up patent expiry, in which expiry advances the laggard only one step (Liu et al. (2020)). With $\zeta \in (-\infty, \infty)$, our specification nests a range of intermediate catch-up speeds. Finally, in an industry with technology gap s, patent expiry occurs at Poisson rate $\eta_s \in [0, \infty)$. We begin our analysis of equilibrium outcomes assuming that IPR policy is uniform: that is, $\eta_s = \eta$ for all $s \in S^+$. Section 8.2 considers state-dependent IPR, in which the patent expiry rate for a given industry depends on its competitiveness.

Entry. We model a fixed mass of potential entrants that cannot produce any intermediate good but can undertake R&D to attempt to enter an industry. Entrants choose an innovation flow rate $x_E \geq 0$ and hire $G_E(x_E) = G(x_E; B_E)$ R&D workers. B_E is an R&D cost scaling parameter that may differ from the scaling parameter for incumbent firms.

In our benchmark model, R&D efforts by potential entrants are *undirected* in the sense of Klette and Kortum (2004) and Akcigit and Kerr (2018). Potential entrants do not know ex ante the product line upon which they will innovate. That is, entrant R&D is not specific to an industry but rather applies to any industry with equal probability. This form of innovation by entrants precludes incumbent leaders from strategically "escaping" the risk of entry—no matter an incumbent's lead, the probability of entry is x_E.[14]

[13] Patent expiry cannot occur, or has no effect, in neck-and-neck industries. That is, $F^p_{0,0} = 1$ and $F^p_{0,\hat{s}} = 0$ for $\hat{s} \neq 1$.

[14] Our model nests quick catch-up models with directed entry (Acemoglu and Cao (2015)): with $\phi_E = \infty$,

At time t, there is a mass 1 of potential entrants that invest in R&D to enter an industry. Upon a successful innovation, the entrant's industry is drawn stochastically from the set of all industries $j \in [0, 1]$. The entrant displaces the industry's follower (or one of the two incumbents, with equal probability, if it enters a neck-and-neck industry). An entrant joining an industry with previous technology gap s begins its life in a random technology position $\hat{s}$ with probability $F^E_{-s,\hat{s}} = F_{-s,\hat{s}}(\phi_E, l_E)$. That is, with higher values of ϕ_E, an entrant is more likely to jump all the way to the maximum leapfrogging position $l_E \geq 0$. If the entrant begins as a leader (with $\hat{s} > 0$), the previous leader becomes the follower. If a potential entrant does not innovate, we assume, as in Akcigit and Ates (2019), that the entrant disappears and is replaced by a new potential entrant.

Recap. Our model features three types of innovation: R&D by incumbent firms at or beyond the reach of laggard leapfrogging (with position $s > l$); R&D by other incumbents ($s < l$); and R&D by potential entrants. These types of innovations affect growth and markups differently. Without leapfrogging, an innovation by tied and leader firms pushes the technology frontier outward but increases the industry's markup. In contrast, innovation by laggards and potential entrants has no direct effect on growth, but reduce the industry's markup. However, innovation by laggards and potential entrants can have an indirect, or composition, effect on *future* growth, if leaders in more competitive industries conduct R&D at higher rates. However, when leapfrogging is possible, innovation by laggards and potential entrants can contribute directly to growth, with ambiguous effects on markups.

2.4 Equilibrium

We focus on Markov perfect equilibria where the strategies of each firm are functions of the payoff-relevant state variables only. We characterize the equilibrium balanced growth path (BGP) in which output and the wage grow at a constant rate g with a stationary distribution of technology gaps. We scale all growing variables by $Y(t)$ so that stationary equilibrium values are constant. The scaled steady-state value function of a firm in position

the BGP equilibrium with undirected entry is the same as the BGP equilibrium with directed entry.

$s \in S$ is given by $v_s(t) \equiv \frac{V_s(t)}{Y(t)}$, where $V_s(t)$ is the firm's discounted expected net profits along the BGP. The labor share is $\omega \equiv \frac{w_t}{Y(t)}$.[15]

Firm and household maximization. Denote the innovation strategy of a firm's competitor by $\{x_s^c\}_{s\in S}$, where x_s^c is the innovation rate of a competitor when the competitor's technology position is s. Each firm chooses its innovation strategy to maximize discounted expected future profits, taking as given the labor share ω, the innovation strategy of its competitor x^c, and the entry rate x_E. The firm trades off R&D costs against the expected discounted profits from an innovation. For a firm in technology position $s \in S$, define the capital gain from a successful innovation Δv_s:

$$\Delta v_s = \sum_{\hat{s}=s+1}^{\bar{s}} F_{s,\hat{s}} v_{\hat{s}} - v_s \tag{5}$$

where for a firm in position s, an innovation will advance the firm to position $\hat{s} \in \{s+1, ..., \bar{s}\}$ with probability $F_{s,\hat{s}}$.

Under preferences (1), the household maximization problem yields the standard Euler equation

$$\frac{\dot{C}}{C} = \frac{\dot{Y}}{Y} = g = r - \rho. \tag{6}$$

A firm's optimal innovation strategy is characterized by the value function v_s. The firm takes into account that operating profits are taxed at rate τ and R&D expenditures are subsidized at rate $\tau_{R\&D}$. For a leader (with $s > 0$), using the Euler equation, the steady state value function satisfies

$$\begin{aligned}\rho v_s = \max_{x_s} \; & \underbrace{(1-\tau)\mathscr{L}_s}_{\text{Operating profits}} - \underbrace{(1-\tau_{R\&D})G(x_s)\omega}_{\text{R\&D costs}} + \underbrace{x_s \Delta v_s}_{\text{Own innovation}} + \\ & \underbrace{x_{-s}^c \sum_{\hat{s}=-\bar{s}}^{s-1} (F_{-s,-\hat{s}} v_{\hat{s}} - v_s)}_{\text{Competitor innovation}} + \underbrace{x_E \sum_{\hat{s}=-\bar{s}}^{s-1} (F_{-s,-\hat{s}}^E v_{\hat{s}} - v_s)}_{\text{Entry}} + \underbrace{\eta_s \sum_{\hat{s}=0}^{s-1} (F_{s,\hat{s}}^p v_{\hat{s}} - v_s)}_{\text{Patent expiry}}.\end{aligned} \tag{7}$$

[15]To ease the notation, we drop the reference to *time* throughout this section.

The leader value function consists of five main parts.[16] The first part is after-tax operating profits, which depend on the Lerner index $\mathscr{L}_s$. The second part is R&D costs and the expected capital from own innovation. The third and fourth parts are the expected loss from competitor innovation and new entrants. Competitor innovation, with arrival rate x^c_{-s}, and entry, with arrival rate x_E, reduce a leader firm's technology position by at least one step. If the competitor innovates, the competitor will advance to position $\hat{s}$ with probability $F_{-s,\hat{s}}$, implying that the firm in position s falls to position $-s$. Similarly, a successful entrant to an industry occupies position $\hat{s}$ with probability $F^E_{-s,\hat{s}}$. If the entrant occupies position $\hat{s}$, then the new position of the incumbent leader is $-\hat{s}$. The final element of the value function reflects patent expiry, which has arrival rate η_s and which reduces a leader firm's technological position by at least one step. Conditional on patent expiry, the probability of the leader moving to position $\hat{s}$ is $F^p_{s,\hat{s}}$.

The value function for a laggard (i.e., a firm with $s < 0$) is

$$\rho v_s = \max_{x_s} -\underbrace{(1-\tau_{R\&D})G(x_s)\omega}_{\text{R\&D costs}} + \underbrace{x_s \Delta v_s}_{\substack{\text{Own}\\ \text{innovation}}} + \underbrace{x^c_{-s}\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{-s,-\hat{s}}v_{\hat{s}} - v_s)}_{\text{Competitor innovation}} - \underbrace{x_E v_s}_{\text{Entry}} + \underbrace{\eta_s \sum_{\hat{s}=0}^{s-1}(F^p_{s,\hat{s}}v_{\hat{s}} - v_s)}_{\text{Patent expiry}}. \tag{8}$$

The value function of a laggard differs from a leader's in two main ways. First, a laggard earns no operating profits. Second, entry displaces the incumbent laggard, whereas the incumbent leader remains a competitor in the industry.

The value function for a tied firm $(s = 0)$ has a form similar to the laggard's, except that a tied firm has an even chance of exiting when an entrant joins an industry

$$\rho v_s = \max_{x_s} -\underbrace{(1-\tau_{R\&D})G(x_s)\omega}_{\text{R\&D costs}} + \underbrace{x_s \Delta v_s}_{\text{Own innovation}} + \underbrace{x^c_{-s}\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{-s,-\hat{s}}v_{\hat{s}} - v_s)}_{\text{Competitor innovation}} + \underbrace{x_E \sum_{\hat{s}=-\bar{s}}^{s-1}(\frac{1}{2}F^E_{-s,-\hat{s}}v_{\hat{s}} - v_s)}_{\text{Entry}}. \tag{9}$$

[16]The appendix derives (7)–(9) from unscaled value functions.

For a firm in any position $s \in S$, profit maximization implies

$$x_s = G'^{-1}(\frac{1}{1-\tau_{R\&D}}\frac{\Delta v_s}{\omega}). \tag{10}$$

The firm innovation rate x_s is therefore increasing in the capital gain from innovation and decreasing in the cost of R&D as measured by the scaled after-subsidy wage $(1-\tau_{R\&D})\omega$.[17]

Denote by $\{\mu_s\}_{s\in S^+}$ the share of industries with a technology gap s. The entrant trades off the cost of R&D against the increased likelihood of successfully entering an industry, with the industry selected at random upon an entrant innovation. For a potential entrant, the expected capital gain from entry is

$$\Delta v_E = \sum_{s=0}^{\bar{s}} \mu_s \left(\sum_{\hat{s}=-s+1}^{\bar{s}} F^E_{-s,\hat{s}} v_{\hat{s}} \right). \tag{11}$$

A successful entrant begins life in an industry with gap s with probability μ_s. Conditional on entering an industry with gap s, the entrant occupies position $\hat{s}$ with probability $F^E_{-s,\hat{s}}$.

The entrant solves:

$$\max_{x_E} -(1-\tau_{R\&D})G(x_E)\omega + x_E \Delta v_E, \tag{12}$$

and the entrant first-order condition is given by:

$$x_E = G_E'^{-1}(\frac{1}{1-\tau_{R\&D}}\frac{\Delta v_E}{\omega}), \tag{13}$$

symmetric with the optimal innovation rate of an incumbent firm (equation (10)).

Market clearing and aggregate growth. With unit inelastic labor supply, labor market clearing requires

$$1 = G_E(x_E) + \sum_{s=0}^{\bar{s}} \mu_s \left[G(x_s) + G(x_{-s}) + (1-\mathscr{L}_s)\frac{1}{\omega} \right]. \tag{14}$$

where $G_E(x_E)$ is potential entrants' demand for R&D workers. In an industry with gap s,

[17] These conclusions are obtained noting that G is twice differentiable and strictly convex. Thus, G'^{-1} is differentiable and strictly increasing.

the demand for R&D workers is $G(x_s) + G(x_{-s})$ and the demand for production labor is inversely related to the Lerner index and the labor share. Because labor is the only factor of production, $(1 - \omega)$ is the share of pure profits in value added.

Our model is a closed economy without (physical) capital, implying that aggregate output equals consumption, $Y = C$. Using the assumptions of a linear production technology for intermediate goods and Cobb-Douglas production of the final good, the wage satisfies

$$w = Q\lambda^{-\sum_{s=0}^{\bar{s}} s\mu_s} \tag{15}$$

and aggregate output satisfies

$$Y = \frac{Q}{\omega}\lambda^{-\sum_{s=0}^{\bar{s}} s\mu_s}, \tag{16}$$

where $Q = \exp(\int_0^1 \ln q_j dj)$ is a "quality index."

In an industry with gap $s \in S^+$, the frontier technology advances in expectation

$$g_s = (\ln\lambda)\left(\sum_{\hat{s}=1}^{\bar{s}}(x_E F^E_{-s,\hat{s}} + x_{-s}F_{-s,\hat{s}})\hat{s} + x_s \sum_{\hat{s}=s+1}^{\bar{s}} F_{s,\hat{s}}(\hat{s} - s)\right) \tag{17}$$

In an industry with gap $s > 0,$ there are two potential sources of frontier-advancing innovation. The first is leapfrogging by the entrant or the incumbent laggard. The arrival rate of entrants in each industry is x_E. Conditional on entry, the frontier advances $\hat{s} > 0$ steps if the entrant leapfrogs the incumbent leader to achieve position $\hat{s}$. Conditional on entry into an industry with technological gap $s,$ the probability of the entrant achieving position $\hat{s}$ is $F^E_{-s,\hat{s}}$. The arrival rate of laggard innovation in an industry with gap s is x_{-s}. Similarly to the case with entry, the innovating laggard can advance the technological frontier by $\hat{s} > 0$ steps if it leapfrogs the incumbent leader to achieve position $\hat{s}$. Conditional on a laggard innovation in an industry with gap $s,$ the probability of advancing to position $\hat{s}$ is $F_{-s,\hat{s}}$. The second source of frontier-advancing innovation is coming from the leader in each industry. A leader innovation arrives at rate x_s in an industry with gap s and advances the leader to position $\hat{s}$ with probability $F_{s,\hat{s}}$, thereby pushing forward the frontier by $\hat{s} - s$

steps. Growth of the quality index Q is therefore:

$$g = \sum_{s=0}^{\bar{s}} \mu_s g_s. \tag{18}$$

Along the balanced growth path, with a stationary distribution of technology gaps and a constant labor share, (18) also provides the growth rate of aggregate output.

The steady-state industry distribution $\{\mu_s\}_{s=0}^{\bar{s}}$ can be obtained by equating outflows and inflows for each gap $s \in (1, ..., \bar{s})$:[18]

$$\begin{aligned}(x_s + x_{-s}(1 - F_{-s,s}) + x_E(1 - F^E_{-s,s}) + \eta_s)\mu_s = \\ \sum_{\sigma \in S^+ \setminus s} \mu_\sigma \left(x_\sigma F_{\sigma,s} + x_{-\sigma}(F_{-\sigma,s} + F_{-\sigma,-s}) + x_E(F^E_{-\sigma,s} + F^E_{-\sigma,-s}) + \eta_\sigma F^p_{\sigma,s} \right)\end{aligned} \tag{19}$$

This identity equates the total outflow from state s (the left hand side) to the total inflow into state s (the right hand side). There are four sources of outflows: innovation by the leader (with arrival rate x_s), innovation by the follower (arrival rate x_{-s}), entry (arrival rate x_E), and patent expiry (arrival rate η_s). The total outflow from state s is the share of gap-s firms times the sum of these arrival rates, taking into account that there is no exit if a laggard jumps to an s step lead (with probability $F_{-s,s}$ conditional on laggard innovation) or an entrant joins with an s step lead (with probability $F^E_{-s,s}$ conditional on entry).

Inflows into state s can occur from an industry with gap $\sigma \in S^+ \setminus s$.[19] A σ-gap industry shifts to an s-gap industry if the leader innovates to obtain position s (with arrival rate $x_\sigma F_{\sigma,s}$) or if a patent expiry closes the gap to s (with arrival rate $\eta_\sigma F^p_{\sigma,s}$). The gap also changes from σ to s if the laggard reaches or the entrant joins at position s or $-s$ (with arrival rate $(x_{-\sigma} + x_E)(F_{-\sigma,s} + F_{-\sigma,-s})$, if the advancement function for laggards and entrants

[18] As noted above, to ease the notation, we intentionally omit the *time* index. The distribution μ_s is uniquely determined by (19), which applies to $s \in \{1, ..., \bar{s}\}$, together with the normalization $\sum_{s \in S^+} \mu_s = 1$. For completeness, the inflow-outflow equation for $s = 0$ is

$$(2x_0 + x_E)\mu_0 = \sum_{\sigma \in S^+ \setminus 0} \mu_\sigma \left(x_\sigma F_{\sigma,0} + x_{-\sigma} F_{-\sigma,0} + \eta_\sigma F^p_{\sigma,0} + x_E F^E_{-\sigma,0} \right).$$

[19] The notation $S^+ \setminus s$ denotes the set S^+ excluding the element s.

is identical with $F = F^E$).

Definition. A balanced growth path equilibrium is, for every t, the tuple $\Upsilon \equiv (g, \omega, \{\mu_s\}_{s\in S^+}, \{x_s\}_{s\in S}, x_E)'$ and $(p_j, y_j)_{j\in[0,1]}$ such that (i) prices and quantities of intermediate goods (p_j, y_j) are consistent with intermediate-good demand and limit pricing; (ii) $\forall s$, x_s is a best response to $\{x^c_\sigma\}_{\sigma\in S}$ and x_E; (iii) $\forall s$, $x^c_s = x_s$ (symmetry); (iv) x_E maximizes entrant expected profits; (v) capital, labor, and goods markets clear; (vi) the distribution of technology gaps $\{\mu_s\}_{s\in S^+}$ is stationary (constant); (vii) growth is determined by (18).

3. Real Interest Rate and the Innovation Multiplier

This section presents our analytical framework for studying the effects of lower interest rates. Section 3.1 studies the effect of a lower discount factor on an individual firm's R&D decision, holding all factors external to the firm constant. We link the effect on an individual firm's decision to the duration of profits from innovation, which varies in the cross section of firms with a firm's position. Section 3.2 introduces our innovation multiplier, mapping the cross section of individual firm R&D effects of a lower discount rate into the total effect on growth. Section 3.3 studies welfare.

3.1 Duration and the individual firm's problem

In this section we study the innovation strategy of a firm, assuming all the conditions external to the firm are fixed: the wage, competitors' strategies, and the discount rate.[20] From the first order condition (10), holding other conditions external to the firm fixed, the effect of an increase in the discount rate on a firm's innovation rate in each state s is

$$\frac{\partial x_s}{\partial \rho} = -c_s \times \mathscr{D}_s \times \Delta v_s, \tag{20}$$

[20] Formally, we define the individual firm's innovation strategy as the R&D strategy satisfying the value functions (7)–(9), taking as given the scaled wage ω, the competitor's strategy $\{x^c_s\}_{s\in S}$, and the entrant strategy x_E. Similarly, the individual entrant's strategy takes as given the wage and the incumbent strategy $\{x_s\}_{s\in S}$.

where $c_s = ((1-\tau_{R\&D})G''(x_s)\omega)^{-1} > 0$ is a curvature term capturing how the marginal cost of innovation changes with the innovation rate, $\mathscr{D}_s \equiv -\frac{\partial \Delta v_s}{\partial \rho}\frac{1}{\Delta v_s}$ is the duration of profits from innovation, and Δv_s is the capital gain from innovation (5).[21]

The curvature term is positive because R&D costs are convex in the innovation rate.[22] Duration, a concept from asset pricing, is the (negative) semi-elasticity of the value of an asset with respect to the discount rate: the percent change in the asset's value from a marginal decline in the discount rate.[23] Duration is measured in time units and captures the amount of time that elapses before an asset holder receives the asset's cash flows.

Equation (20) shows that the sign and magnitude of the effect of a lower discount rate on an individual firm's R&D depends on the *time pattern* of the expected profits from innovation. With respect to the sign, if the expected profits from innovation are weakly positive at all horizons $z > t$, then their duration must be positive. That is, holding factors external to the firm constant, a lower discount rate is associated with higher R&D. However, if the expected profits from innovation are negative at some future dates $z > t$, then the discount-rate valuation effect can in principle be negative even as the capital gain from innovation is positive. This outcome can obtain when an innovation leads to positive profits shortly after the innovation, followed by losses later.[24] With a negative duration of profits from innovation, an individual firm facing a lower discount rate would decrease R&D.

3.2 The innovation multiplier

The effects of a change in the household discount rate are, ultimately, a general equilibrium (GE) question. The previous analysis of firm-level innovation abstracts from changes in the wage and in competitor and entrant innovation strategies. Moreover, aggregate vari-

[21]Similarly, the effect of a higher discount rate on entrant R&D, holding factors external to the entrant constant, satisfies to first order, $\frac{\partial x_E}{\partial \rho} = -c_e \times$ (Duration of profits from entry) $\times \Delta v_E$.

[22]The curvature term c_s is increasing in the innovation rate x_s if and only if $\gamma \leq 0.5$. In the benchmark model, we set $\gamma = 0.5$, implying that the curvature term is invariant across s.

[23]The Appendix provides a formal definition of the expected profits from innovation. The duration of expected profits from innovation is connected to, but distinct from, the widely cited metric of the duration of firm profits, or, $\mathbb{D}_s \equiv -\frac{\partial v_s}{\partial \rho}\frac{1}{v_s}$. See the Appendix for a mapping from the duration of firm profits to the duration of profits from innovation.

[24]In our model, an innovation can lead to losses at some horizons due to an endogenous increase in R&D and therefore R&D costs.

ables such as growth and the average markup depend not only on the changes in firm-level innovation rates, but also on the resulting changes in composition of technology gaps.

To inspect these forces, we provide a mapping from the cross section of individual firm R&D effects of a lower discount rate (holding constant all factors external to the firm) into the GE effect on growth, the distribution of technology gaps (and hence markups), and firm R&D. Denote the cross section of individual firm R&D effects effects of a lower discount by $\boldsymbol{\partial x} = \begin{bmatrix} \partial x_{-\bar{s}} & \dots & \partial x_{\bar{s}} & \partial x_E \end{bmatrix}'$. Further, recall that the BGP elements are grouped in the vector Υ (Section 2.4). Then, for the broad family of models considered in this paper, the following result holds.

Theorem 1 ***(The Innovation Multiplier)***. *Consider a change in the discount rate ρ. There exists a general equilibrium multiplier matrix $\mathbb{M}$ such that the effect of this change on the BGP, $\Upsilon \equiv (g, \omega, \{\mu_s\}_{s \in S^+}, \{x_s\}_{s \in S}, x_E)'$, to a first order, is*

$$\boldsymbol{d\Upsilon} = \mathbb{M}\, \boldsymbol{\partial x} \tag{21}$$

where $\boldsymbol{\partial x}$ is the cross section of individual firm R&D effects (expression (20)).

Corollary 1 *The multiplier $\mathbb{M}$ for changes in policy parameters—the mapping from their individual firm effects on R&D into their GE effects on growth and other outcomes—does not depend on the particular policy change driving these firm level effects.*

Corollary 1 highlights that the same multiplier, $\mathbb{M}$, can be used to assess the GE effects of a broad range of policy changes that affect the cross section of firm R&D decisions. Such policy changes include changes to the R&D subsidy ($\tau_{R\&D}$), the corporate tax rate (τ), the patent expiry rate (η), or targeted changes in these policies (gap-dependent changes such as changes affecting only laggards or entrants).

The difference between firm-level R&D effects and GE outcomes arises because general equilibrium imposes additional conditions: labor market clearing, the stationarity of the technology gap distribution, and the strategic requirement that $\boldsymbol{\partial x}$ is a best response. Individual firm effects perturb each of these general equilibrium conditions. By knowing how each condition is perturbed, one recovers the mapping from the individual effects to the change in GE outcomes.[25]

The innovation multiplier $\mathbb{M}$ maps $\boldsymbol{\partial x}$ to the general equilibrium effect on the entire equilibrium vector Υ. Each row of $\mathbb{M}$ maps $\boldsymbol{\partial x}$ into the GE effect on a particular element on Υ. As a matter of notation, we call the row corresponding to growth $\mathbb{M}_g$, the row corresponding to the share of industries with gap 3 $\mathbb{M}_{\mu_3}$, and so on. We refer to the element of the multiplier matrix that maps ∂x_s into growth by $\mathbb{M}_{\partial x_s \to g}$.

Recall that the GE effects of a change in the discount rate $d\rho$ on incumbent strategy in position s, entrant strategy, and the wage are denoted by dx_s, dx_E, and $d\omega$, respectively.

Corollary 2 ***(Decomposition of the Multiplier)*** *To first order, the GE effect on x_s (dx_s) of a change in the discount rate is*

$$dx_s = \underbrace{\frac{\partial x_s}{\partial \rho} d\rho}_{\text{Valuation}} + \underbrace{\frac{\partial x_s}{\partial \omega} d\omega}_{\text{Wage}} + \underbrace{\sum_{\sigma=-\bar{s}}^{\bar{s}} \frac{\partial x_s}{\partial x^c_\sigma} dx^c_\sigma}_{\text{Strategic: Incumbent}} + \underbrace{\frac{\partial x_s}{\partial x_E} dx_E}_{\text{Strategic: Entrant}}, \tag{22}$$

where

$$\frac{\partial x_s}{\partial \omega} = c_s \left(\underbrace{-\frac{\Delta v_s}{\omega}}_{\text{Static}} + \underbrace{\frac{\partial \Delta v_s}{\partial \omega}}_{\text{Dynamic}} \right) \tag{23}$$

is the effect on a firm's innovation rate of a change in the wage (taking the competitor and entrant strategy as given),

$$\frac{\partial x_s}{\partial x^c_\sigma} = c_s \frac{\partial \Delta v_s}{\partial x^c_\sigma} \tag{24}$$

is the effect of a change in a competitor's strategy in position $\sigma \in S$ (taking the wage, entrant strategy, and competitor strategies in other positions as given), and

$$\frac{\partial x_s}{\partial x^c_\sigma} = c_s \frac{\partial \Delta v_s}{\partial x_E} \tag{25}$$

[25] Our mapping from firm individual R&D effects to GE effects resembles that of Auclert and Rognlie (2018) for the *consumption multiplier*. Their paper studies the household saving decision with idiosyncratic income risk and borrowing constraints. Our analysis differs in an important way: Strategic and labor-market interactions imply that the cross section of individual firms R&D effects is required to calculate GE outcomes. That is, the innovation multiplier is a vector. In contrast, in their setting, the GE change in consumption is a scalar multiplier times the sum of partial equilibrium effects.

is the effect of a change in the entrant's strategy.

Corollary 2 thus unpacks the GE effect on a firm's innovation strategy into different components: valuation, wage, and strategic, respectively. (Later, Figures 9 and 11 illustrate these components.)

In addition, a higher wage affects a firm's innovation rate through two channels: it raises the cost of R&D today (the static component) and it affects the capital gain from an innovation through its effects on future R&D costs (the dynamic component). This dynamic component is:

$$\frac{\partial \Delta v_{s,t}}{\partial \omega} = \left(\sum_{\hat{s}=s+1}^{\bar{s}} F_{s,\hat{s}} \mathbb{E}_{\hat{s},t} - \mathbb{E}_{s,t} \right) [\int_t^\infty -e^{-\rho z}(1-\tau_{R\&D})(\frac{x_\varsigma}{B})^{\frac{1}{\gamma}})dz]. \tag{26}$$

An increase in the competitor's innovation rate when the competitor is in position σ affects a firm's innovation rate x_s in a similar way, by changing the probability of a competitor innovation today if $s = -\sigma$ and by changing the capital gain from innovation.

The terms $\frac{\partial \Delta v_s}{\partial \omega}$ and $\frac{\partial \Delta v_s}{\partial x_\sigma^c}$ are inputs into the innovation multiplier $\mathbb{M}$. These terms enter into $\mathbb{M}$ in a complex way, reflecting the interaction of a rich set of general equilibrium forces. The Appendix shows that $\frac{\partial \Delta v_s}{\partial \omega}$ and $\frac{\partial \Delta v_s}{\partial x_\sigma^c}$ can be obtained as a function of the firm value function $\boldsymbol{v}$; the firm innovation strategy $\boldsymbol{x}$; the patent expiry rate η; and the parameters governing the speed of catch-up following innovation ϕ and patent expiry ζ.

There is an analogue of corollary 2 for the partial equilibrium change in the entry rate x_E,

$$dx_E = c_e \left(\frac{\partial \Delta v_E}{\partial \rho} d\rho + \left(\frac{\partial \Delta v_E}{\partial \omega} - \frac{\Delta v_E}{\omega} \right) d\omega + \Sigma_{\sigma=-\bar{s}}^{\bar{s}} \frac{\partial \Delta v_E}{\partial x_\sigma^c} dx_\sigma^c + \frac{\partial \Delta v_E}{\partial x_E} dx_E \right). \tag{27}$$

3.3 Welfare, markups, and interest rates

In the absence of patent policies—or, more broadly, other methods of restricting competitors' access to innovation (Hall et al. (2014))—firms have no incentive to innovate and the balanced growth path equilibrium features zero growth and zero markups. Patents can partly align the social and private benefits of innovation by providing market power to innovators, but this market power distorts production decisions. A further social cost of patents is that innovation decisions of laggards and potential entrants are motivated partly the prospect of innovations that do not advance the frontier but are privately beneficial by increasing the likelihood of future leadership. The patent system therefore gives rise to a rich set of tradeoffs (Arrow (1962), Nordhaus (1969), and Acemoglu and Akcigit (2012), among others). We explore these tradeoffs in our model, with a focus on how welfare and optimal policies change with the discount rate. We focus on two key innovation policies: the (possibly gap-dependent) patent expiry rate $\{\eta_s\}_{s\in S^+}$ and the R&D subsidy $\tau_{R\&D}$.

In a balanced growth path equilibrium, the consumption of the representative household is

$$C(t) = Y(t) = Y(0)e^{-gt}. \tag{28}$$

Innovation policies therefore have a level effect and a growth effect on the consumption path. The level effect is captured by $Y(0)$ and reflects that policies affect market power and thereby the static equilibrium allocation studied in Section 2.2. An increase in production distortions due to market power shifts the path of consumption downward multiplicatively. The growth effect reflects how patent policies affect the cross section of firm innovation strategies. An increase in growth tilts the path of consumption. In our benchmark quantitative analysis, we find that shifts in policy that increase innovation in partial equilibrium lead to higher growth but also higher markup distortions. Remarkably, for some alternative parametrizations, such shifts in policy lead to higher growth and lower markup distortions—pointing again to the general equilibrium interactions highlighted by the innovation multiplier.

With log preferences, along a BGP equilibrium, welfare depends in a simple way on the initial level of output and and growth. Welfare at time $t = 0$ is $\mathbb{W}(0) = \int_{t\in[0,\infty)} e^{-\rho t} \ln(C(t))\, dt$, or equivalently

$$\mathbb{W}(0) = \frac{\ln(Y(0))}{\rho} + \frac{g}{\rho^2} \tag{29}$$

$$= \frac{1}{\rho} \underbrace{\ln Q(0)}_{\text{Initial condition}} + \frac{1}{\rho} \underbrace{\left(-\ln\omega + \sum_{s\in S^+} \mu_s \ln(1-\mathscr{L}_s) \right)}_{\text{Ratio of output to potential output}} + \frac{g}{\rho^2}, \tag{30}$$

where the ratio of output to potential corresponds to $\ln \frac{Y(0)}{Q(0)}$. The welfare measure is increasing in growth and the initial level of output; the relative weight on growth rises as the discount rate declines. The difference between output $Y(0)$ and potential output $Q(0)$ reflects how market power distorts production. Conditional on the labor share, the difference between output and potential output depends on both the level and dispersion of the Lerner index. The difference increases in response to higher or more disperse Lerner indexes, in the sense of first and second order stochastic dominance, respectively. Market power decreases the aggregate demand for production labor by a factor $\sum_{s\in S^+} \mu_s(1-\mathscr{L}_s)$, relative to aggregate demand absent market power.[26] Note that this effect of markups on the aggregate demand for labor is homogenous of degree 1 in the cross section of Lerner indexes, $\{\mathscr{L}_j\}_{j\in[0,1]}$. When the cross section of Lerner indexes is non-degenerate, market power also gives rise to misallocation across industries, in the form of differences in revenue productivity. The equilibrium static allocation (the price and quantity of each intermediate good) is identical to what would obtain in an environment with perfect competition and free access to the leading production technology in each industry, but with a revenue tax equal to $\mathscr{L}_s$.

Later, we analyze how patent policy and R&D subsidies affect welfare, and how optimal policy changes with the discount rate.[27] A change in the discount rate will lead, endogenously, to a change in the "frontier" of growth rates g and market-power distortions $Y(0)-Q(0)$ that the social planner can choose from. The change in optimal policy will reflect this change in the social planner's frontier, as well as the change in how the social planner trades off higher growth and lower market-power distortions, as reflected in equation (30).

4. Calibration

[26]We say an economy is without market power when the Lerner index $\mathscr{L}_s$ is equal to 0 for all j.
[27]As in Acemoglu and Akcigit (2012), we maximize balanced growth path welfare conditional on $Q(0)$.

4.1 Moments and identification

The model has 14 structural parameters. We fix the discount rate ρ equal to 2 percent which roughly corresponds to an annual discount factor of 98 percent. We also fix the curvature of the R&D cost function (i.e., the elasticity of successful innovation with respect to R&D), $\gamma = 0.5$, consistent with empirical estimates discussed in Acemoglu et al. (2018). The corporate tax rate τ is 20 percent and the R&D subsidy $\tau_{R\&D}$ is also 20 percent, as in Akcigit and Ates (2019).[28]

There are 9 remaining structural parameters to be calibrated. We identify these parameters using an indirect inference approach (Lentz and Mortensen (2008) and Akcigit and Kerr (2018)). We obtain some of the model moments directly from our solution for the balanced growth path equilibrium Υ. For the remaining moments, we compute model values using a simulation of $N = 50,000$ industries for $T = 12$ years.[29] We compare these moments in the model and the data, choosing the parameters such that they minimize the criterion

$$\min \sum_{m=1}^{14} weight_m \left(\frac{\text{model}(m) - \text{data}(m)}{\frac{1}{2}|\text{model}(m)| + \frac{1}{2}|\text{data}(m)|} \right)^2. \tag{31}$$

Following Acemoglu et al. (2018), we give aggregate growth and the average markup a weight five times the weight of the other moments.

We choose aggregate and cross-sectional moments for identification based on key features emphasized in the Schumpeterian growth literature.

Productivity growth. We target a productivity growth rate of 1.6 percent, representing an average of the measured total factor productivity growth rate for 1960–2017, from Fernald et al. (2017), and the growth rate of innovative firms in Acemoglu et al. (2018).

Lerner index. From Hall (2018), we target a mean Lerner index of 0.145 and an across-industry standard deviation of 0.11. We also target the implied net markup of 0.194.

Firm-level innovative output. Kogan et al. (2017) estimate the economic value of inno-

[28] Note that the value of the R&D subsidy does not affect the equilibrium BGP produced by the calibration because changes the subsidy are equivalent to changes in the R&D cost scale parameters B and B_E, for a given value of the parameter γ.

[29] Note that we solve the model at a monthly frequency and simulate the model at the frequency that is monthly or higher. A higher frequency is needed for the calibration of some parameter values, so that the probability of no change in a given industry in any period remains positive.

vations based on the stock market reaction to patent grants. Kogan et al. (2017) define the innovation output of a firm in a given year as the sum of the economic value of all patents earned, normalized by firm value. We target the mean and the 50th and 90th percentiles of firm-level annual innovation output. These moments are obtained in the data for the sample period 1950–2010 used by Kogan et al. (2017), restricting attention to firms with positive R&D. We calculate these moments in the model using the simulation.

Growth decomposition. Foster et al. (2001) develop a simple growth decomposition using the identity,

$$\Delta\Theta_t = \underbrace{\sum_{i\in C_t} \kappa_{it-1}\Delta\theta_{it}}_{\text{within}} + \underbrace{\sum_{i\in C_t}(\theta_{it-1}-\Theta_{t-1})\Delta\kappa_{it}}_{\text{between}} + \underbrace{\sum_{i\in C_t}\Delta\theta_{it}\Delta\kappa_{it}}_{\text{cross}} + \underbrace{\sum_{i\in N_t}\kappa_{it}(\theta_{it}-\Theta_{t-1})}_{\text{entry}} + \underbrace{\sum_{i\in X_t}\kappa_{it-1}(\Theta_{t-1}-\theta_{it-1})}_{\text{exit}}, \quad (32)$$

where C_t is the set of continuing firms, N_t of entering firms, and X_t of exiting firms, in industry j between $t-1$ and t, respectively. In addition, $\theta_{it} = \ln(\frac{y_{it}}{l_{it}})$ is the (log) productivity of firm i at time t, $\kappa_{it} = \frac{l_i t}{\sum_i l_i t}$ is the activity share of firm i at time t, and $\Theta_{it} = \sum_i \kappa_{it}\theta_{it}$. This decomposition separates the change in industry productivity growth into five components: within-firm productivity improvements; the between-firm effect; the cross effect; the entry effect; and the exit effect. We target the within and entry shares from Foster et al. (2008), with the remaining components studied, but untargeted.

Profit volatility. We calculate the standard deviation of firm-level growth in operating profits between year y and $y+1$. Following Kogan et al. (2017) and others, we calculate operating profits for publicly traded firms as sales (SALE) minus cost of goods sold (COGS) using COMPUSTAT. We estimate an unconditional standard deviation of 47.8 percent. Conditional on having profits in the top profit quintile in year y, the standard deviation of profit growth is estimated to be 20.3 percent. Values for the innovation-output moments and the profit volatility moments, from both the model and the data, are obtained after winsorizing at the 1st and 99th percentiles. We restrict attention to firms with positive R&D.

	Parameter	Description	Value
1.	λ	Innovation step size	1.02
2.	ϕ	Catch-up speed, laggard innovation	-0.13
3.	ϕ_E	Catch-up speed, entrant innovation	2.51
4.	B	Scale parameter, incumbents' R&D cost function	2.93
5.	B_E	Scale parameter, entrants' R&D cost function	0.23
6.	l	Laggard leapfrogging (number of steps)	0.00
7.	l_E	Entrant leapfrogging (number of steps)	3.00
8.	ζ	Catch-up speed, patent expiry	2.90
9.	η	Patent expiry rate	0.05

Table 1: **Fitted parameter values for the baseline economy with entry.**

Young firm employment share. Following Decker et al. (2014), we obtain employment share by firm age from the US Census's Business Dynamics Statics. We target the employment share of young firms, defined as less than or equal to 10 years old, of 26 percent, the average for the earliest and most recent years with publicly available cross-sectional data (1987 and 2014). To further gauge the role of entry, we calculate, in the same way, the employment share of firms less than 1 year old, which is 3.9 percent in the data.

R&D intensity. From Acemoglu et al. (2018), the ratio of R&D to sales, conditional on a firm being small and young, is 6.4 percent. The ratio conditional on being large and old is 3.7 percent.[30]

4.2 Calibrated parameters and equilibrium properties

Table 1 reports calibrated parameter values. For the incumbent innovation catch-up parameter, we obtain $\phi = -0.13$. This value implies that laggard advancement, conditional on innovation, is very close to the slow catch-up assumption of Liu et al. (2020). Our calibrated value for the incumbent leapfrogging parameter, $l = 0$, implies no leapfrogging, also consistent with a slow catch-up laggard innovation. In contrast, the calibrated catch-up parameter for patent expiry is $\zeta = 2.25$, implying a high degree of quick catch-up. For a laggard in an industry with the median technology gap ($s = 8$), the probability of advanc-

[30]Following Acemoglu et al. (2018), we calculate conditional R&D intensity in the model defining young firms as being less than or equal to 9 years old and small firms as having firm size below the median.

Moment	Model	Data	Source
1. Growth rate	1.58	1.60	AAABK, Fernald
2. Lerner index: Mean	14.90	14.50	Hall
3. Lerner index: Standard deviation	10.20	11.00	Hall
4. Markup: Mean	19.50	19.40	Hall
5. Innovation output: Mean	5.88	6.38	Kogan
6. Innovation output: 50th pctile	0.00	0.00	Kogan
7. Innovation output: 90th pctile	15.50	17.90	Kogan
8. FHK "Within" contribution to growth	73.30	67.10	FHS
9. FHK "Entry" contribution to growth	24.60	24.00	FHS
10. Young firm employment share	27.90	26.20	BDS
11. R&D intensity: Young & small	7.09	6.40	AAABK
12. R&D intensity: Old & large	3.40	3.70	AAABK
13. Profit volatility: Unconditional	41.20	47.80	COMPUSTAT
14. Profit volatility: Top quintile	22.30	20.30	COMPUSTAT

Table 2: **Model and data moments.** Sources: AAABK: Acemoglu et al. (2018). Fernald: Fernald et al. (2017). Hall: Hall (2018). Kogan: Kogan et al. (2017). FHS: Foster et al. (2008). BDS: US Census Business Dynamics Statistics, Public Release.

ing more than one step following an innovation is extremely low (0.2 percent) but much higher following a patent expiry (above 99 percent).

The entrant catch-up parameter is $\phi_E = 2.51$ implying that entrants tend to move to a much more advantageous position than an innovating laggard. Moreover, entrants can leap ahead of the incumbent leader by as many as $l_E = 3$ steps. Entrants therefore lead to an economy with much more dynamism than would obtain with slow catch-up laggard innovation and without entry.

There is a very good fit between our model-implied moments and the data. The model is able to capture relevant features of the data, including the aggregate growth rate, average markup and Lerner index, and the cross section of profit volatility and R&D intensity. The model matches well the entire distribution of innovation output, including the high share of firms with no innovation output in a given year and the right tail of firms with quite high innovation output (left panel of Figure 3). There is also a good fit between the Foster et al. (2001) growth decomposition in the model and the data (middle panel of Figure 3). Finally, the model matches well the distribution of markups estimated by Hall (2018), as shown in the right panel of Figure 3.

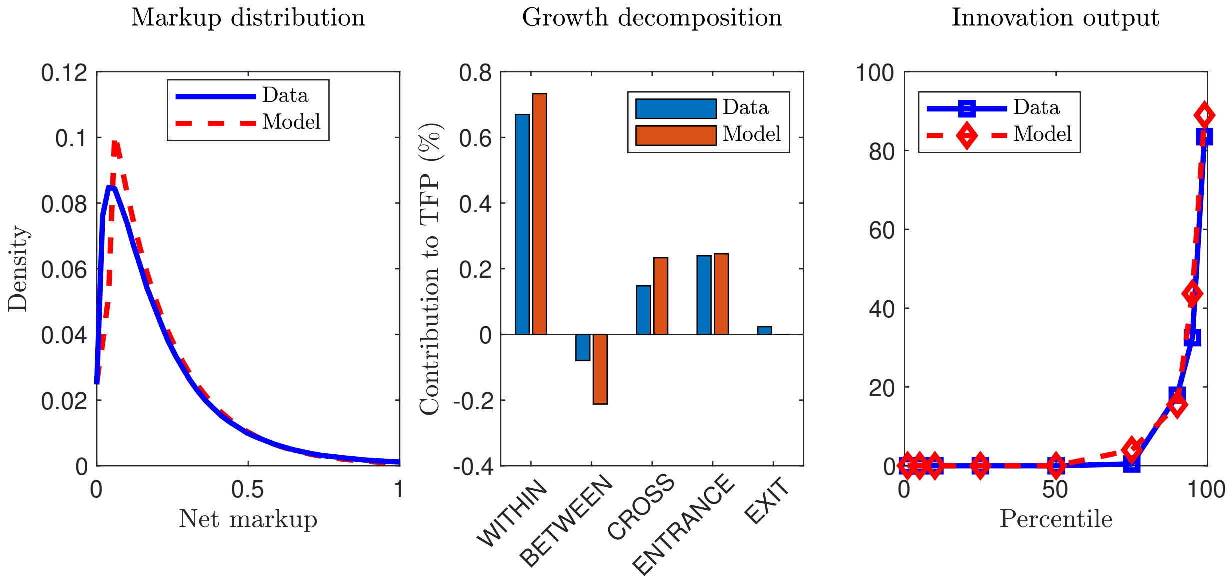

Figure 3: **Innovation output, growth decomposition, and markup distribution.** The left panel shows the distribution of markups, from Hall (2018). The middle panel shows the Foster, Haltiwanger, and Krizan growth decomposition, using data from Foster et al. (2008). The right panel shows the distribution of innovation output, from Kogan et al. (2017).

The model delivers a 1.6 percent annual equilibrium growth rate, driven by innovations of incumbent leaders and entering firms. Figure 4 shows the innovation strategy as a function of a firm's technology position. R&D expenditures are high for tied firms —the *escape competition* effect of Aghion et al. (2001), reflecting the winner-takes-all nature of Bertrand competition. Because of the slow catch-up nature of laggard innovation, laggards more than a few steps behind "give up" and choose very low innovation rates. Farther-behind laggards are discouraged from innovating because they can only reach leadership and positive operating profits through a series of successive innovations unmet by leader innovations. The low innovation rates of far-behind laggards have an important strategic effect on the innovation of slightly-ahead leaders, which innovate at very high rates in order to expand their lead, which would allow them to face a discouraged laggard. Once leaders attain an advantage of more than a few steps, they face a discouraged laggard and, absent entry, can maintain or expand their lead while reducing their R&D expenditures. As their lead builds further, these firms reduce their R&D expenditures progressively, reflecting

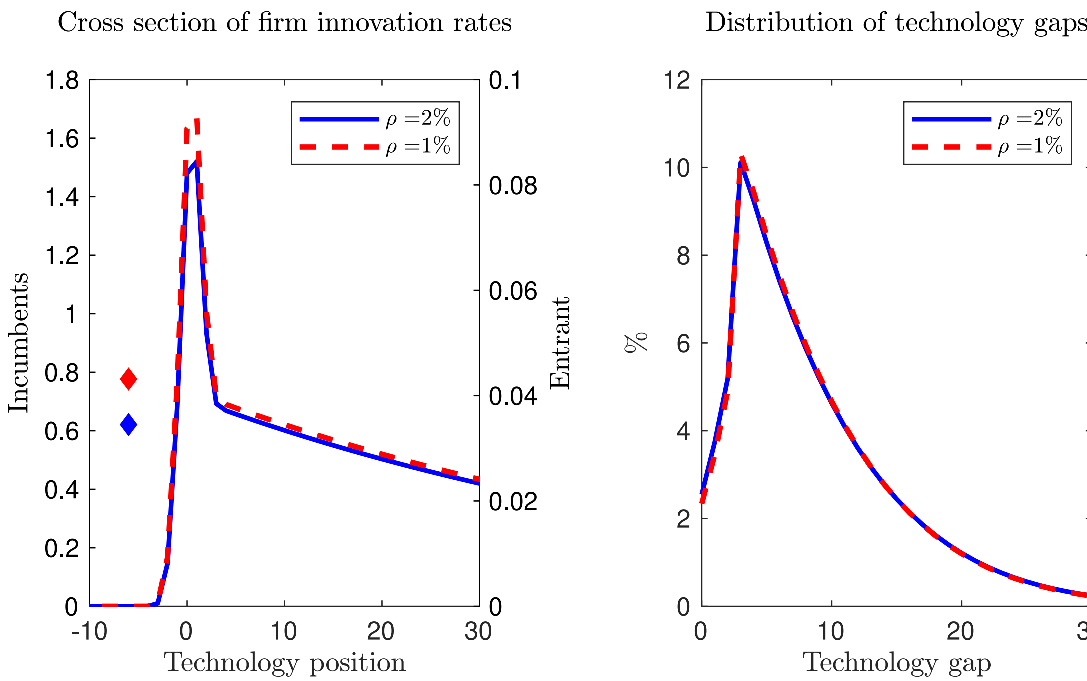

Figure 4: **Firm innovation strategy and the distribution of technology gaps.** The continuous lines in the left panel show the annualized innovation rates of incumbents, as a function of their technology position, for two different values of the discount rate ρ. The diamonds show the annualized innovation rates of entrants. The right panel shows the distribution of technology gaps, $\{\mu_s\}_{s\in S^+}$.

that operating profits for leaders are concave in their advantage. Thus, leader innovation is higher in more competitive industries. This pattern implies that, all else equal, growth is higher in economies with smaller technology gaps.

The innovation rates of incumbent leaders imply, on average, between 0.5 and 1.5 innovations per year, depending on the leader's position. The innovation rate for potential entrants is much lower, implying an entry rate of 3.5 percent per year. This entry rate generates an employment share of firms less than one year old equal to 3.7 percent, close to the value in the data. Even with this low entry rate, entrants have an important role in our economy, because of their quick catch-up speed (high ϕ) and leapfrogging ($l_E > 0$). As a result, entrants contribute to growth by moving the frontier if they leapfrog and by increasing the share of high-innovation, competitive industries. Through the lens of the FHK growth decomposition, entrants account for 25 percent of productivity growth, in line with the data.

Our benchmark model features undirected entry. As noted in Section 2.3, our model nests directed-entry models with quick catch-up. Because our calibrated model features a quite high speed of entrant catch-up, if we study a version with directed entry, the quantitative BGP equilibrium will be almost identical to the equilibrium with undirected entry (as shown in the Internet Appendix).

The strategic and composition effects described in this section are key to understand the transmission of changes in the household discount rate throughout the economy, which we turn to next.

5. Growth, markups, and the discount rate

This section studies how the discount rate affects growth, innovation, and the distribution of markups. Section 5.1 addresses this question using our calibrated model, which matches well cross sectional moments related to markups, profit volatility, R&D, and innovation output. Section 5.2 assess robustness, by altering assumptions related to entry, labor supply, and the elasticity of substitution across intermediate goods. Section 5.3 studies an alternative economy in which the forces of creative destruction are severely curtailed: in addition to having no entry, laggard innovation is subject to "pure" slow catch-up.

5.1 Benchmark economy

Figure 5 shows the two key relations that determine growth and the interest rate. The blue line, which we call the **innovation schedule**, captures the interplay between innovation, patent policy, and competition on the firm side of the economy. This mapping from the real interest rate to the growth rate is consistent with profit maximization by firms and aggregate firm labor demand equal to the inelastic labor supply. For the calibration that matches key economic moments, productivity growth rises monotonically and approximately linearly as the interest rate falls. A lower interest rate boosts growth mostly along the intensive margin. As shown in Figure 4, as the discount rate falls, firm innovation strategies shift upward. There is also a complex extensive margin, or composition, effect. As the interest rate declines, the equilibrium features a lower share of industries with very large technology gaps. This pro-competitive shift contributes positively to growth, because innovation decreases in the technology gap. However, as the interest rate declines, there are also fewer industries with very tight technology gaps, which contributes negatively to growth. That is, mass shifts from very low and very high gaps toward intermediate gaps. Overall, a decline in the discount rate induces a change in composition that slightly boosts growth.

The dashed lines in Figure 5 represent the **Euler equation**, or the household side of the economy, for different values of the household discount rate ρ. Conditional on the discount rate ρ, the equilibrium growth and interest rate are determined so that, given that interest rate, firms' innovation decisions lead to output growth of g, while households increase their consumption at the same rate g. Therefore, as the discount rate ρ declines, equilibrium growth increases and the real rate falls. Because growth increases as the discount rate declines, the real interest rate declines less than one-for-one with the household discount rate. Thus, the innovation side of the economy serves effectively to dampen the transmission of discount rate changes into changes in the equilibrium real rate.

The average markup varies little with the interest rate (Figure 5, middle panel). The full distribution of markups is also little affected by a change in the interest rate (right panel). Consistent with the discussion of the composition effect on growth, as the interest rate declines, there are slightly smaller shares of industries with very low or very high markups, which on balance explains the small decline in average markups.

Elasticity of intertemporal substitution (EIS). The appendix shows how the model can be extended to allow an EIS not equal to 1, at some computational cost. As the EIS varies, the innovation schedule from Figure 5 is unaffected—because the EIS affects only households' consumption-saving decision. A lower EIS rotates the Euler equation to be flatter, raising the interest rate and lowering growth. Regardless of the EIS, a decline in the discount rate leads to a rise in growth. The relation between growth and the interest rate traced out as the discount rate varies is unaffected by the EIS.

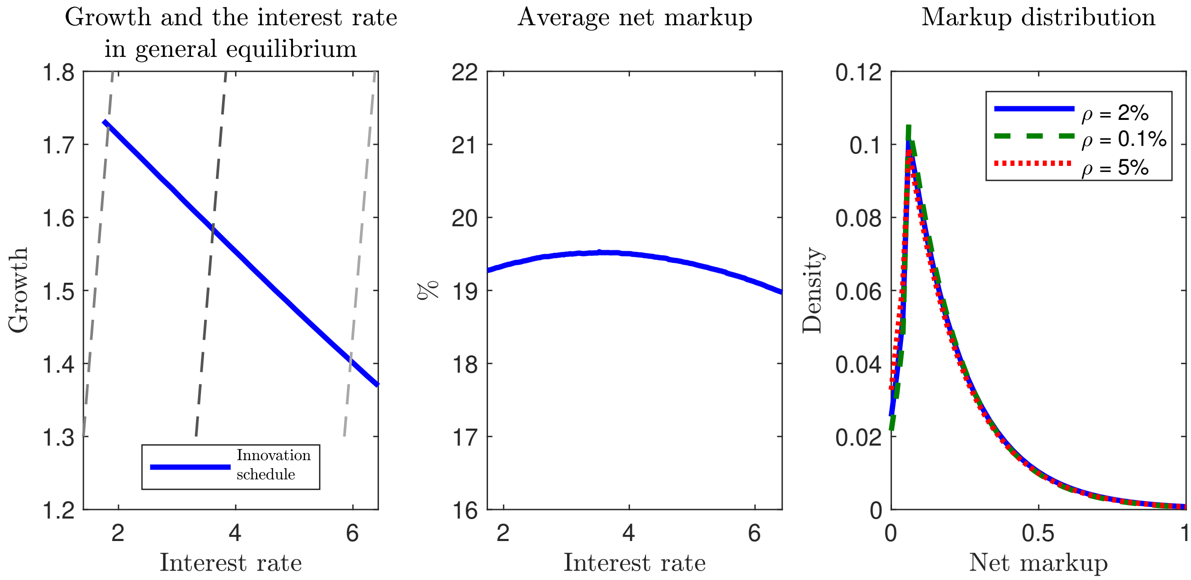

Figure 5: **Growth, markups, and the interest rate.** Gray dashed lines in the left panel show the Euler equation for ρ equal to 0.1, 2, and 4.5 percent.

5.2 Robustness

In the results so far, entry has an important role, consistent with the shares of employment and growth accounted for by entrants and young firms. This section assesses whether key properties of the BGP equilibrium, including the monotonic negative relation between growth and the interest rate, depend on the inclusion of entrants. To do so, we turn entry "off" and recalibrate the parameters. We also assess whether our main results depend on our assumptions of inelastic labor supply and within-industry perfect substitution between intermediate goods.

No entry. Table 3, column (1), shows the results from calibrating a version of the benchmark model without entry. The recalibration targets the non-entry moments from the original calibration. Were entry to be turned off without recalibrating the parameters, the distribution of markups would shift to the right (with a mean of 40 percent) and become more disperse (with a standard deviation of 36 percent), rendering the model unable to match the data. Thus, when we recalibrate, to hit targets related to market power and profits, the speed of laggard innovation (parametrized by ϕ) rises dramatically relative to the economy with entry. Creative destruction shifts from entrants to the laggards. With this recalibration, non-entry moments are matched fairly well despite some mismatch with the growth decomposition due to the absence of entry. In this no-entry economy, growth rises as the interest rate declines (Figure 6, top left panel, solid blue line). A sharp rise in laggard innovation as the interest rate falls reflects a high duration of laggard profits from innovation and implies a decline in the average markup (top right panel).

Elastic labor supply. The models studied thus far feature inelastic labor supply, as in most canonical Schumpeterian models. Table 3, column (2), shows the results from calibrating the no-entry economy, with perfectly elastic labor supply.[31] The calibrated parameters and model moments are very similar to the model with inelastic labor supply. Further, growth monotonically declines with the interest rate (Figure 6, top left panel, dashed pink line). This finding that the elasticity of labor supply has almost no effect on key properties of the equilibrium is, of course, conditional on the parameters in this exercise. Below, when pure slow catch-up in laggard innovation is imposed as an assumption, the elasticity of labor supply is quite important.

Imperfect substitution among within-industry intermediate goods. For Table 3, column (3), we drop the assumption that, within a given industry $j \in [0, 1]$, the two competing firms produce perfect substitutes. We assume instead a constant elasticity of substitution $\kappa > 1$, fixing $\kappa = 12$ as in Liu et al. (2020). The quality of fit with the non-entry moments is moderate and worse than when we assumed perfect substitution ($\kappa = \infty$). The relation between growth and the interest rate remains negative and monotonic (Figure 6, top left panel, dotted green line).

5.3 The growth "speed limit" economy

Next, we study a version of the model that severely restrict the forces of creative destruction, by assumption. We assume that there is no entry and impose extremely slow innovation catch-up ($\phi = -\infty$). We set the patent catch-up parameter ζ so that the probability of a quick catch-up conditional on patent expiry is very small (0.8%).[32] Recalibrating all the

[31] The Appendix characterizes the BGP equilibrium with perfectly elastic labor supply.
[32] With $\zeta = -\infty$, we obtain a degenerate economy with zero growth regardless of the discount rate.

Model Features	(1)	(2)	(3)	(4)	Data
Pure slow catch-up innovation	No	No	No	Yes	
Labor supply	Inelastic	Elastic	Elastic	Inelastic	
Within-industry elasticity of substitution	Perfect	Perfect	Imperfect	Perfect	
Parameters					
Innovation step size (λ)	1.04	1.04	1.10	1.03	
Catch-up speed, laggard innovation (ϕ)	0.36	0.38	1.08	$-\infty$	
Scale parameter (B)	1.40	1.44	0.86	1.89	
Laggard leapfrogging, number of steps (l)	0.00	0.00	0.00	0.00	
Catch-up speed, patent expiry (ζ)	2.68	2.68	-1.31	0.00	
Patent expiry rate (η)	0.00	0.00	0.00	0.24	
Within-industry elasticity of substitution (κ)	∞	∞	12	∞	
Moments					
Growth rate	1.65	1.67	1.60	1.64	1.60
Lerner index: Mean	15.00	14.60	16.80	38.10	14.50
Lerner index: Standard deviation	9.02	9.12	2.00	12.50	11.00
Markup: Mean	19.10	18.60	20.30	68.40	19.40
Innovation output: Mean	5.66	5.63	1.34	0.65	6.38
Innovation output: 50th pctile	0.00	0.00	0.00	0.00	0.00
Innovation output: 90th pctile	16.40	17.70	8.13	2.40	17.90
FHK "Within" contribution to growth	87.10	86.90	75.80	98.80	67.10
Profit volatility: Unconditional	39.40	39.00	21.50	10.40	47.80
Profit volatility: Top quintile	20.30	20.20	16.80	7.41	20.30

Table 3: **Model and data moments for the no-entry economies.** Column (1) shows parameters and model moments when we assume no entry and recalibrate the economy. Column (2) shows results when perfectly elastic labor supply is assumed. Column (3) shows results when further assuming that within-industry intermediate goods are imperfect substitutes, with constant elasticity of substitution κ. Column (4) shows the speed-limit economy.

other parameters conditional on these assumptions, we find that a very high patent expiry rate implies a fit with non-entry moments that is only moderately worse. Interestingly, in this slow catch-up economy, there is no growth "speed limit." Growth rises monotonically as the interest rate falls, reaching a plateau at very low interest rates (Figure 6, bottom left panel, solid pink line).

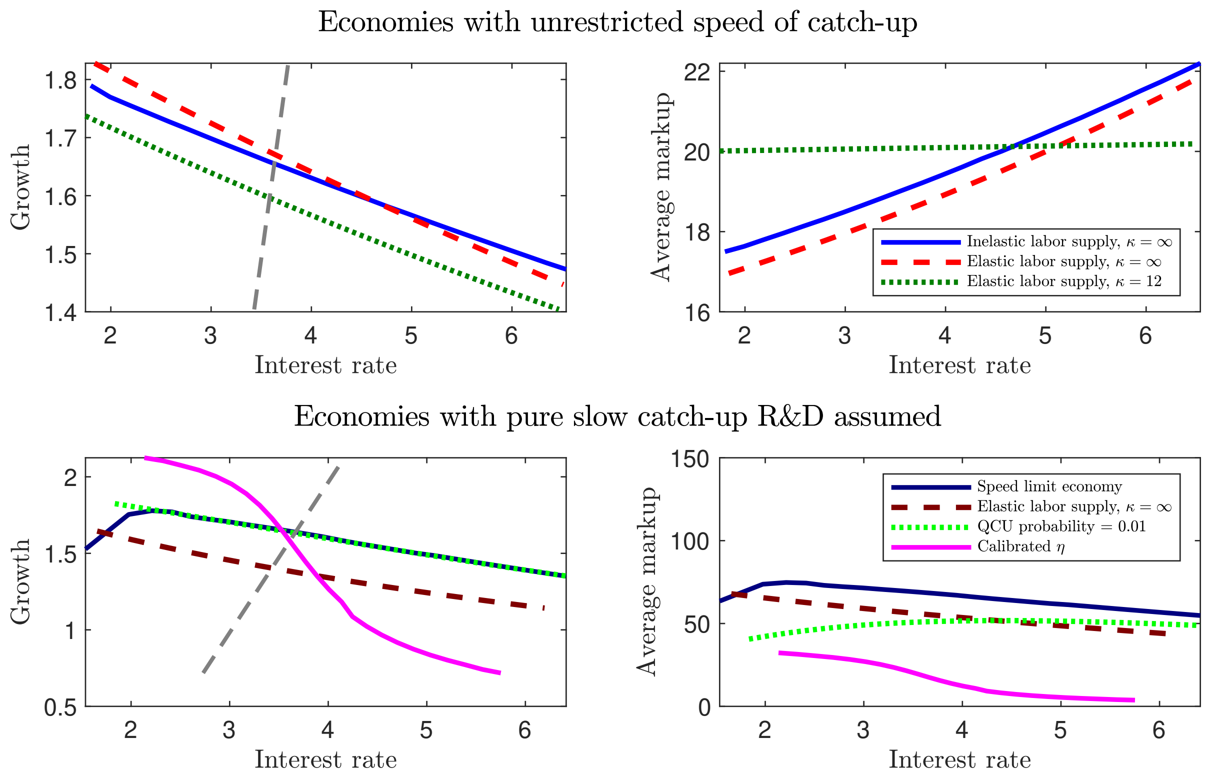

Figure 6: **Growth, markups, and the interest rate in no-entry economies.** The top panels pertain to economies in which the speed of catch-up is unrestricted and calibrated to match the data, as shown in Table 3. The bottom panels pertain to economies in which laggard innovation is restricted to advance the laggard only one step (pure slow catch-up). For each economy, the relation of growth and the average markup with the interest rate is traced out for $\rho \in [0.01, 5]$.

To further curtail creative destruction, we then lower the patent expiry rate from its calibrated annual value (54%) to 24%. Table 3 shows the resulting parameter values and moments. The average markup (near 70 percent) and the standard deviation of markups (almost double the value from Hall (2018)) are counterfactually high.

This speed-limit economy is incredibly sclerotic. For example, profit volatility is much too low relative to the data. Correspondingly, the distribution of innovation output is extremely concentrated at low levels (middle panel of Figure 7)—in stark contrast with the data and the benchmark model. Moreover, the FHK "within" contribution to growth reaches almost 100 percent, with reallocation playing almost no role.

This sclerotic economy features a growth "speed limit," with an inverted-U relation

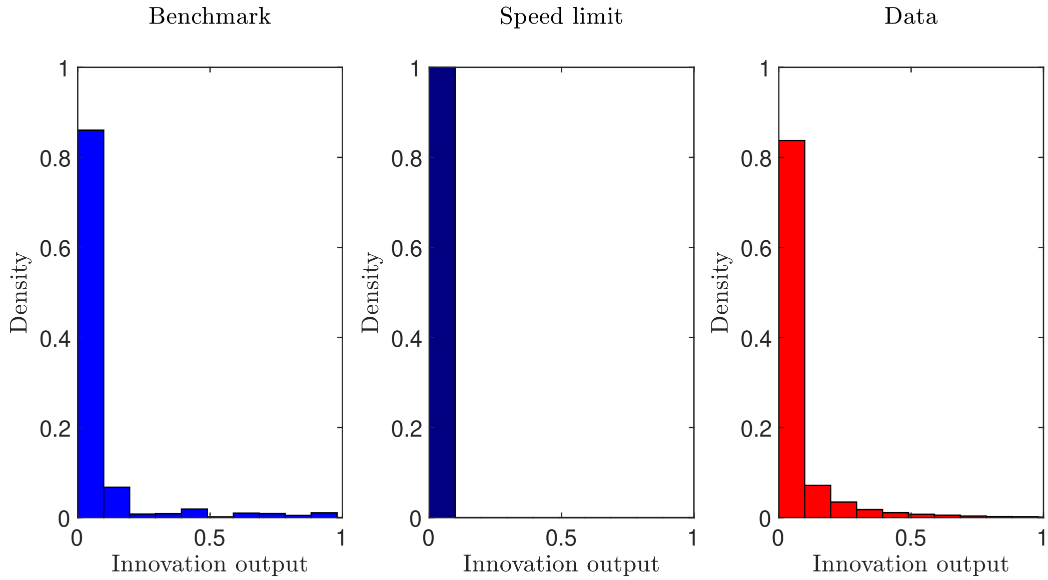

Figure 7: **Distribution of innovation output.** The figure shows the densities of innovation output in the benchmark economy (left panel), the speed limit economy (middle panel), and the data (right panel). Innovation output is calculated as in Kogan et al. (2017), confining attention to firms with positive R&D.

between growth and the interest rate (Figure 6, bottom left panel, solid blue line). Let's explore the mechanism. In competitive industries, a lower interest rate boosts R&D (Figure 8). Starting from a high interest rate, growth rises, as the intensive margin (higher innovation rates) dominates the extensive margin or composition effect (lower growth due to more industries with low-innovation, high-markup superstars). However, as the interest rate falls further, higher innovation by prospective superstars boosts the wage, in part because these high-innovation firms have low R&D productivity (i.e., G is convex). The higher wage crowds out innovation by entrenched superstars. That is, as shown in Figure 8, starting from a very low interest rate, a marginally lower interest rate reduces R&D by far-ahead leaders.[33] At low interest rates, far-ahead leaders predominate in the sclerotic economy. Therefore, a lower interest rate, by crowding out their R&D, leads to *lower* aggregate growth.

[33]The decline in innovation rates of far-ahead leaders is self-reinforcing, because it leads to a pro-competition composition effect. This composition effect raises labor demand because more competitive industries, with smaller markup distortions, employ more production workers.

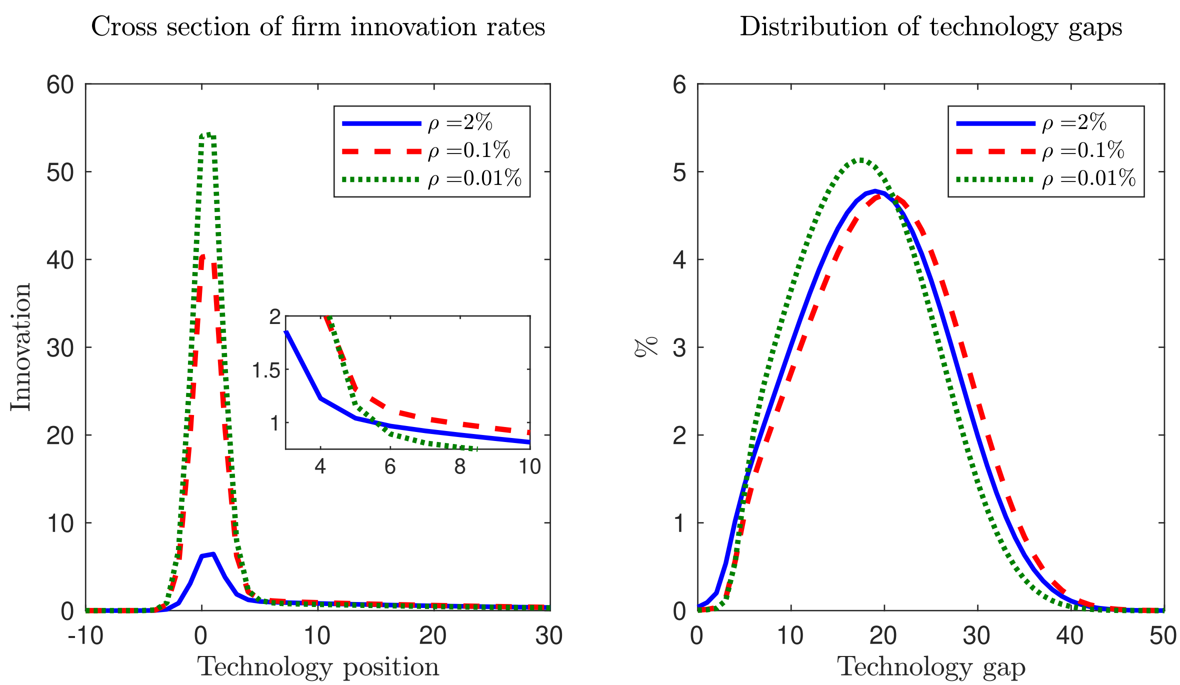

Figure 8: **Firm innovation strategy and the distribution of technology gaps in the speed limit economy.** The inset box of the left panel shows that, in the speed limit economy described in column (4) of Table 3, a decline in the discount rate from 0.1% to 0.01% leads to lower innovation for far-ahead leaders.

The labor market is crucial to our growth speed limit. If we assume a perfectly elastic labor supply, the speed limit disappears, with growth rising monotonically as the interest rate falls (Figure 6, bottom left panel, dashed brown line). The inverted-U relation between growth and the interest rate is fragile in other ways as well. With only some departure from the limiting case of "pure" slow catch-up, the inverted-U relation vanishes (bottom left panel, dotted green line).[34] Moreover, as we will explore in the next section, adjusting the uniform patent expiry rate optimally eliminates the speed limit.

5.4 The technology of R&D and the elasticity of labor supply

This section reviews how the technology of R&D and the elasticity of labor supply affect the equilibrium relations between growth, markups, and the interest rate. Our benchmark model, calibrated to match cross sectional moments related to innovation, competition, and creative destruction, features *intermediate* speeds of catch up through laggard innovation, patent expiry, and entrant innovation. That is, a laggard or entrant innovation or a patent expiry implies a positive probability of closing the technology gap completely. These parameters (and, relatedly, the patent expiry rate η) were chosen through indirect inference. In this benchmark economy, growth rises as the interest rate falls (Section 5.1). This result is robust to turning off entry, making labor supply elastic, and assuming imperfect substitution among competing goods within an industry (Section 5.2). In the benchmark economy, the distribution of markups varies little with the interest rate. Without entry, the average markup declines as the interest rate falls.

Our speed-limit economy, in contrast, features pure slow catch-up from laggard innovation (Section 5.3. That is, laggard innovation closes the productivity gap only stepwise. For the speed-limit economy, we do not assume pure slow catch-up through patent expiry. Instead, we assumed a small probability of quick catch-up through patent expiry (0.8%). In this economy, we obtain an inverted-U relation between growth and the interest rate (and also an inverted-U relation between the average markup and the interest rate). However, the speed limit economy does not match the data well and the growth speed limit is

[34] The economy with some departure from pure slow catch-up sets ϕ such that there is a 1% probability of quick catch-up, conditional on a laggard innovation.

extremely fragile. With elastic labor supply, the speed limit disappears, because it arose from prospective superstars crowding out R&D by far-ahead leaders. The growth speed limit also disappears with moderate deviations from pure slow catch-up in laggard innovation or if the patent expiry rate is set to match aggregate growth and the average markup.

Remark. Liu et al. (2020) investigates the relation of growth, market power, and the interest rate under alternative assumptions. Their model assumes no entry and that laggard innovation and patent expiry close the technology gap only stepwise. In their theoretical analysis, an inverted-U relation between growth and the interest rate always obtains.

Labor supply is inelastic in our main exercises, but perfectly elastic in their model. We have three main conclusions with respect to labor supply. First, when we calibrate our model to match cross sectional moments, there is no inverted-U in the growth-interest rate relation, whether labor supply is elastic or not. Second, in our speed-limit economy, the inverted-U vanishes if we assume elastic labor supply. Their inverted-U, in contrast, is derived specifically with elastic labor supply. This contrast reflects that our inverted-U arises through a different mechanism, driven by different assumptions about the technology of R&D. Third, in our speed-limit economy, when we impose pure slow catch-up in laggard innovation, markups decline as the interest rate falls toward zero, but in their model, markups rise. This difference is rooted in the labor market. In their setting, with perfectly elastic labor supply, the accumulation of market power builds unabated as the interest rate falls. (This unabated accumulation of market power explains why they obtain a growth speed limit: leader innovation declines as the leader gets further ahead, implying that a less competitive grows more slowly, all else equal.) In contrast, in our model with pure slow catch-up through innovation and inelastic labor supply, GE adjustments in the wage dampen and eventually reverse the accumulation of market power.[35]

[35] In the Internet Appendix, we offer a detailed description of the quantitative approach in Liu et al. (2020) and how it further curtails creative destruction, beyond the assumption of pure slow catch-up.

6. Inspecting the mechanism

We will proceed as follows. Section 6.1 studies the effects of a lower discount rate on firm R&D strategy, holding constant all other factors external to the firm. Using the innovation multiplier, Section 6.2 studies the implications for growth. The multiplier unpacks how strategic and labor-market interactions and composition effects dampen, but do not overturn, the effects on firm R&D strategy holding external factors constant.

6.1 The individual firm's response

This section focuses on the response of individual firms to a lower discount rate, taking as given conditions external to the firm—the wage and competitor strategies. For a firm in a given technology position, this effect on firm innovation is the product of the duration of profits from innovation, the expected capital gain conditional on innovation, and a curvature term arising from the innovation production function (equation (20)). Each of these factors will, in general, differ for firms in different technology positions.[36]

Holding all other factors constant, a decline in the discount rate implies a rise in innovation for all firms, in every technology position (grey dashed line in Figure 9) and including entrants (grey diamond). As explained in Section 3.1, the sign of this partial effect, in principle, could be negative if an innovation from a certain technology position leads to back-loaded losses through increased R&D expenditures. The uniformly positive partial effect of a lower discount rate on innovation is therefore a result of the model calibration. In other words, for the parameter values that best match our key moments, the duration of profits from innovation is uniformly positive.

The partial effects of a lower discount rate are largest in tied industries. This result reflects two forces. First, *tied firms' capital gain from innovation* is high. An innovating tied firm starts to earn operating profits and, as its lead increases further, reduces its R&D expenditures. Second, *tied firms' duration of profits from innovation* is higher than leaders'

[36]Because the elasticity of the innovation rate to R&D expenditures, γ, is equal to 0.5, the curvature term is the same for all incumbent firms in our calibration. However, because the R&D scaling parameter for incumbents, B, differs from the scaling parameter for potential entrants, B_E, the curvature terms differ for incumbents and entrants.

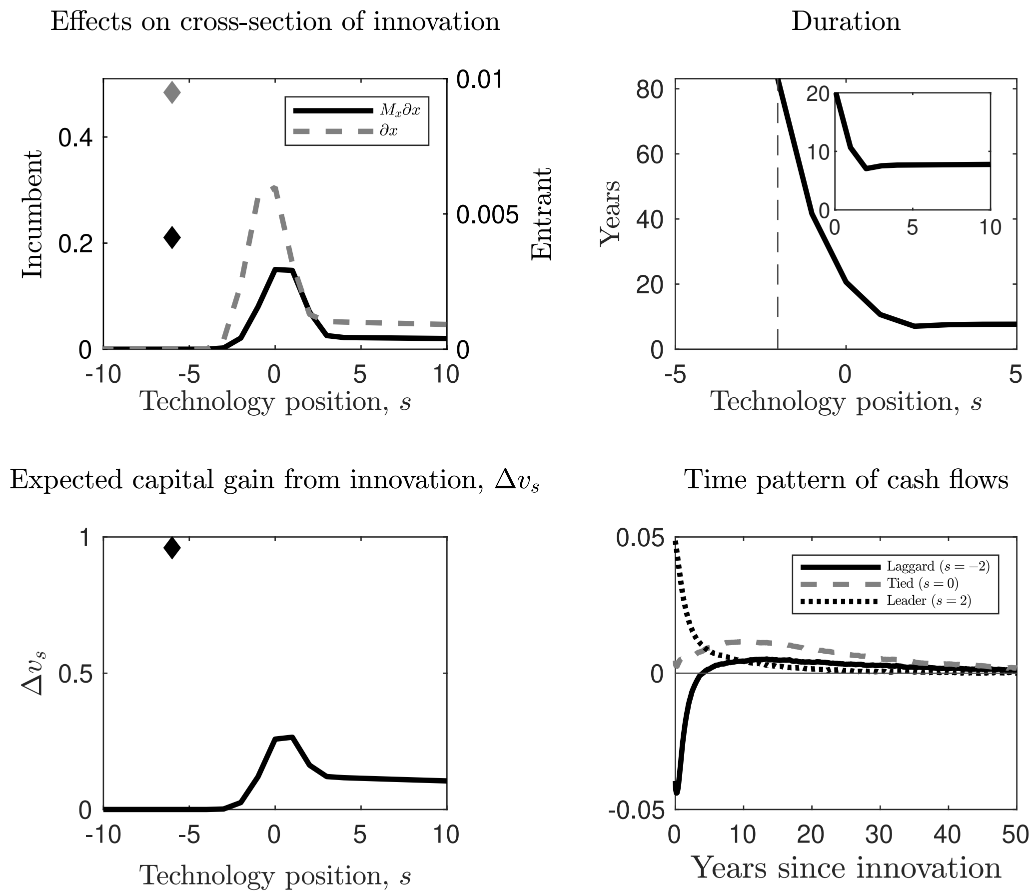

Figure 9: **Effects on firm innovation of a lower discount rate.** The grey dashed line in the top left panel shows the effect of a 100 bps decline in the discount rate, holding all other factors constant (∂x). The solid line shows the general equilibrium effect ($dx = M_x \partial x$). Values for the entrant are shown with a diamond. The capital gain from innovation (bottom left) and the duration of profits from innovation (top right) are key determinants of ∂x (equation (20)). The bottom right panel shows the time pattern of the expected profits from innovation. (Note: For technology gaps $s \leq -3$, the duration is undefined, with the innovation rate and capital gain from innovation positive but very close to 0.)

because tied firms' profits from innovation are relatively backloaded. Upon innovating, a tied firm initially incurs high R&D expenses, to try to build a lead sufficient to discourage its competitor. Thus, for tied firms, the expected profits from innovation accrue steadily over a long period (bottom right panel of Figure 9, dashed line). In contrast, an innovating leader even a couple of steps ahead immediately sees an increase in net profits, but this increase is small, reflecting the concavity of leaders' operating profits in their technology advantage and a small decline in R&D expenses. These leaders' advantage is eroded over time by entrants and patent expiry, making these leaders' profits from innovation relatively front-loaded (dotted line). Compared with nearly tied leaders (in competitive sectors), further-ahead leaders (in less competitive sectors) therefore have a lower capital gain from innovation and a shorter duration of expected profits from innovation. Thus, the partial effect on leaders' innovation declines progressively with leaders' technology advantage.

Interestingly, laggards have much lower capital gains from innovation than tied firms but the duration of their profits from innovation is very high (bottom left panel of Figure 9). This high duration reflects that, for laggards, an innovation is associated with immediately higher R&D expenditures and no change in operating profits. Thus, for laggards, an innovation on impact generates expected losses followed by an increase in expected profits further in the future (bottom right panel of Figure 9, solid line). As laggards fall further behind, their capital gain from innovation declines quickly (faster than the duration of profits from innovation rises). With fairly slow catch-up innovation, laggards in uncompetitive industries have little prospect of closing the technology gap, explaining why their capital gain from innovation—and therefore their innovation rate—is approximately zero.

The effect on entrant innovation (of a 100 basis point decline in the discount rate, holding other factors constant) has a small magnitude in absolute terms. Due to entrant quick catch-up and leapfrogging, the effect on entrant innovation nevertheless will translate into large effects on growth, as will be discussed later. The percent increase in entrant innovation (28 percent) is similar in magnitude to the percent increase in tied-firm innovation (21 percent). These increases in entrant and tied-firm innovation are significantly dampened in general equilibrium.

6.2 The innovation multiplier

The effect of a lower discount rate on growth and innovation is ultimately a general equilibrium question. That is, firms take into account not only the direct valuation effect of a lower discount rate, but also resulting changes in their competitors' strategies and the wage. In addition, because firms in different competitive positions innovate at different

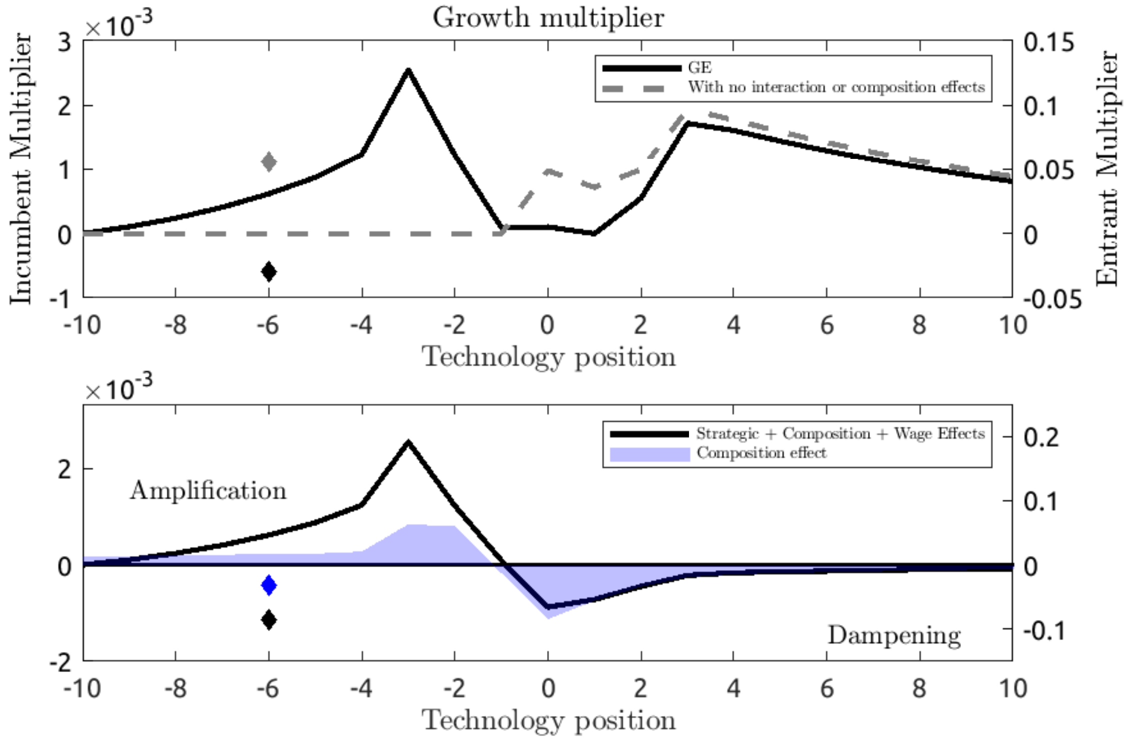

Figure 10: **The innovation multiplier for growth** Top panel: The black line shows the general equilibrium multiplier mapping the cross section of PE effects on R&D into the GE effect on growth. The grey line shows the (fictitious) multiplier that would obtain absent GE forces. Values for the entrant are shown by diamonds. Bottom panel: The black line shows the difference between the GE multiplier and the multiplier that would obtain absent GE forces. The shaded area shows the share of this difference due to composition effects, as defined in footnote 38.

rates, endogenous changes in the composition of technology gaps can boost or diminish aggregate growth.

Theorem 1 shows how the cross section of individual-firm effects on firm R&D ($\boldsymbol{\partial x}$) translate into the effect on aggregate growth ($dg = M_g \boldsymbol{\partial x}$). Consider how ∂x_s, the effect of a lower discount rate on the innovation rate of a firm in position $s \in S$ holding all other factors constant, translates into the effect on aggregate growth. This partial effect ∂x_s potentially contributes to the general-equilibrium aggregate growth effect through five channels. First, for tied and leader firms ($s \geq 0$), innovations contribute directly to growth, by pushing out the technology frontier. However, because there is no leapfrogging, laggard innovations do not contribtue directly to growth. Second, higher innovation by a firm in position s sharpens its competitors' incentive to escape competition through R&D. That is, holding all other factors constant, an increase in innovation by a firm in position s leads firms in position $-s$ to increase innovation: $\frac{\partial x_{-s}}{\partial x_s^c} > 0$.[37] A third, countervailing force arises from *strategic trickle-down* effects. For example, a leader currently 2 steps ahead has less to gain from innovation, the greater the innovation rates of laggards 3 or more steps behind, with whom the leader would compete if the leader were to innovate. That is, holding all other factors constant, $\frac{\partial x_\sigma}{\partial x_s^c} \leq 0$, for $\sigma > -s$. A fourth way that firm-level innovation affects growth is through composition effects. A partial increase in innovation by a firm in position s translates into a change (given by $M_{\partial x_s \to d\mu_\sigma}$) in the share of industries with gap $\sigma \in S^+$.[38] This change in the the economy's competitiveness affects growth because firms in different technology positions innovate at different rates. The fifth channel is that a partial increase in innovation by a firm in position s translates into a change in the wage, given by $M_{\partial x_s \to d\omega}$. This increase in the wage, in turn, affects the innovation rates of all firms, with the effect on innovation for a firm in position σ given by $\frac{\partial x_\sigma}{\partial \omega}$.

The innovation multiplier for growth folds in *all* of the channels, presenting a compact view of strategic, labor market, and composition effects. The top panel of Figure 10 shows the growth multiplier M_g (solid black line). Absent strategic, wage, and composition effects, increases in laggard innovation do not contribute to aggregate growth, as shown by the dashed grey line. However, taking these effects into account, higher laggard innovation in competitive industries boosts growth, mostly by boosting leader innovation (through an escape-competition channel) and a composition effect (a more competitive economy). For a firm in position s, the composition effect on growth is $\sum_{\sigma \in S^+} \mathbb{M}_{\partial x_s \to d\mu_\sigma} g_\sigma$ (shaded region). Interestingly, partial increases in innovation by tied firms and slightly ahead leaders do not contribute to growth, even though their innovations directly advance the technology frontier. This result obtains due to a composition effect: innovation by these firms leads to a less competitive economy. For far-ahead leaders, the growth multiplier is almost exactly equal to what would obtain absent strategic, composition, and wage effects.

[37]See the Appendix for a fuller characterization.

[38] The composition effect on growth of a unit partial increase in x_s is $\sum_{\sigma \in S^+} \mathbb{M}_{\partial x_s \to d\mu_\sigma} g_\sigma$.

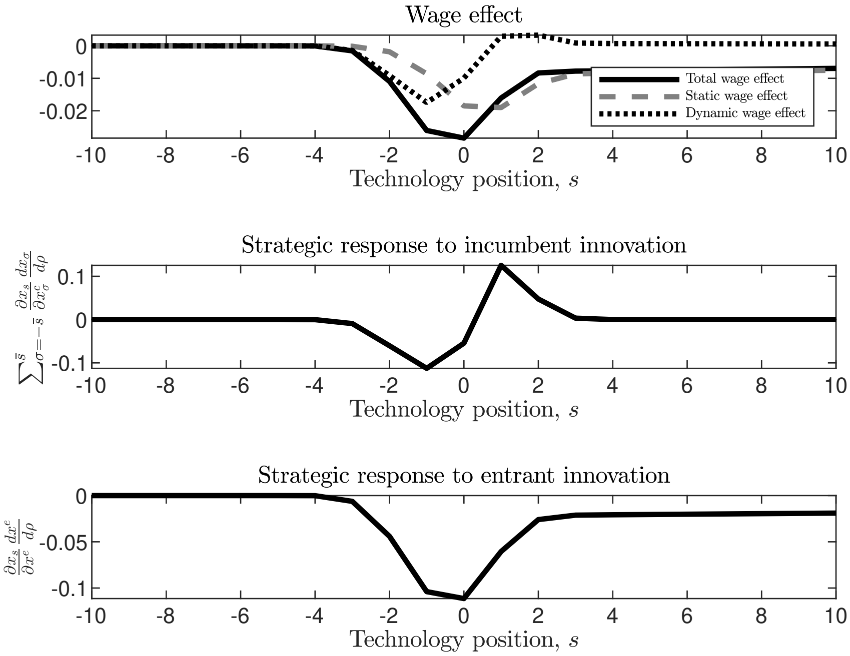

Figure 11: **Wage and strategic general equilibrium effects on the cross section of incumbent firm innovation rates.** Corollary 2 shows that, for a firm in any position s, the difference between the PE and GE effects on innovation can be decomposed into a wage effect (top panel), the strategic effect of changes in incumbent innovation (middle panel), and the strategic effect of the change in entrant innovation (bottom panel). The figure shows this decomposition, for a 100 basis point decline in the discount rate.

Finally, although entrant innovations directly contribute to growth, partial increases in entrant innovation *reduce* aggregate growth, due to a composition effect (due to leapfrogging and quick catch-up by entrants, entry contributes to having fewer highly competitive industries) and a strategic trickle-down effect. That is, a partial increase in entry dissuades leaders from innovation, because such leaders know that the gains from a successful innovation are more likely to be snatched away by an entrant and are thus shorter-lived.

The effects of general equilibrium interactions through the labor market are not readily apparent in Figure 10, but wage effects are in fact quantitatively relevant and vary according a firm's technology position. Corollary 2 decomposes the general equilibrium effect on firm-level innovation rates into three components: valuation, labor market, and strategic. A lower discount rate raises innovation, increasing the demand for scientific labor, with a small indirect effect on the demand for production labor due to minor changes in the distribution of markups. As a result, the wage rises slightly, dampening the rise in innovation (top panel of Figure 11). This wage effect is the sum of a static and dynamic component (equation (23)). The static component, as shown in the corollary, is proportional to the capital gain from innovation. The dynamic component has opposite sign for laggards and leaders. A higher wage increases a leader's wage bill, but by a smaller magnitude for further-ahead leaders (which use less R&D and production labor). Through this channel, a higher wage increases a leader's incentive to innovate, implying a positive dynamic wage component. In contrast, a higher wage bill raises a laggard's wage bill, more so for more competitive laggards. Thus, for laggards, the dynamic wage component is negative.

The partial increase in incumbent innovation induced by a lower discount rate has important strategic effects (middle panel of Figure 11). The strategic effect varies in sign and magnitude with a firm's technology position. In competitive industries, the strategic effect increases leader innovation: higher innovation by laggards sharpens the escape competition motive of leaders in these industries. In contrast, in these industries, the strategic effect reduces laggard innovation: higher competitor innovation makes leadership less durable, discouraging innovation by laggards through a trickle-down effect. Similarly, higher innovation by entrants has a negative strategic effect on all incumbents, through a similar trickle-down effect (bottom panel of Figure 11).

Taken together, the innovation multiplier for growth (Figure 10) and the decomposition of firm-level innovation effects (Figure 11) depict how strategic interactions, the general-equilibrium adjustment in the wage, and composition effects variously amplify and dampen the effect on aggregate growth of a lower discount rate. Holding all other factors constant, a lower discount rate increases innovation and aggregate growth. On balance, these strategic, labor market, and composition forces dampen, but do not overturn, the increase in growth from a lower discount rate.

7. Financial Frictions

This section investigates how access to credit affects the relation of growth, markups, and the interest rate. Financial frictions might be expected to limit creative destruction and hence might foster conditions for an inverted-U relation between growth and the interest rate. Following Aghion et al. (2019a), we assume that firms need to finance (on an intra-period basis) their entire R&D wage bill and that this financing is limited to a fixed share of firm value v_s.[39] Specifically,

$$\omega G(x_s)(1-\tau_{R\&D}) \leq \alpha v_s. \tag{33}$$

This constraint implies that the innovation rate for a firm in position s is

$$x_s = \min\{G'^{-1}(\frac{\Delta v_s}{\omega(1-\tau_{RD})}), G^{-1}(\frac{\alpha v_s}{\omega(1-\tau_{RD})})\}. \tag{34}$$

Per Aghion et al. (2019a), lower values of α capture reduced credit access.

Figure **??**, Panel A shows firm-level innovation strategy for three values of the limited-commitment parameter, for a discount rate of 2%. The constraint binds, if at all, for tied firms and leaders in competitive industries. In contrast, far-ahead leaders are not constrained, because their desired innovation rate is low and their value is high.[40] The constraint thus inhibits a firm's initial accumulation of market power but not the maintenance of a large advantage. Accordingly, for lower values of $\alpha,$ the distribution of gaps is more competitive (Panel B). Reduced credit access leads to lower growth, because escape-competition R&D in competitive industries is a key driver of growth in the benchmark economy, and lower markups (Panels C and D).

In the speedlimit economy, reduced credit access also constrains tied firms and leaders in more competitive industries, but the implications for growth and markups are very different. At high interest rates, reduced credit access reduces growth and markups, as in the benchmark economy. However, at low interest rates, reduced credit access can *increase* growth, by shutting down the growth-stifling composition effects of lower interest rates in the speed-limit economy (Panel G). Moreover, when credit access is most restricted, the

[39]Another example of firms having to finance their working capital on an intra-period basis is Jermann and Quadrini (2012). The limited commitment constraint in our model can be micro-founded by assuming that, in the event of default, firms can renege and make a take-it-or-leave-it offer to creditors, which can only seize a fraction of the value of the firm.

[40]Laggards are not constrained, because with fairly slow catch-up, their desired innovation rate is very low.

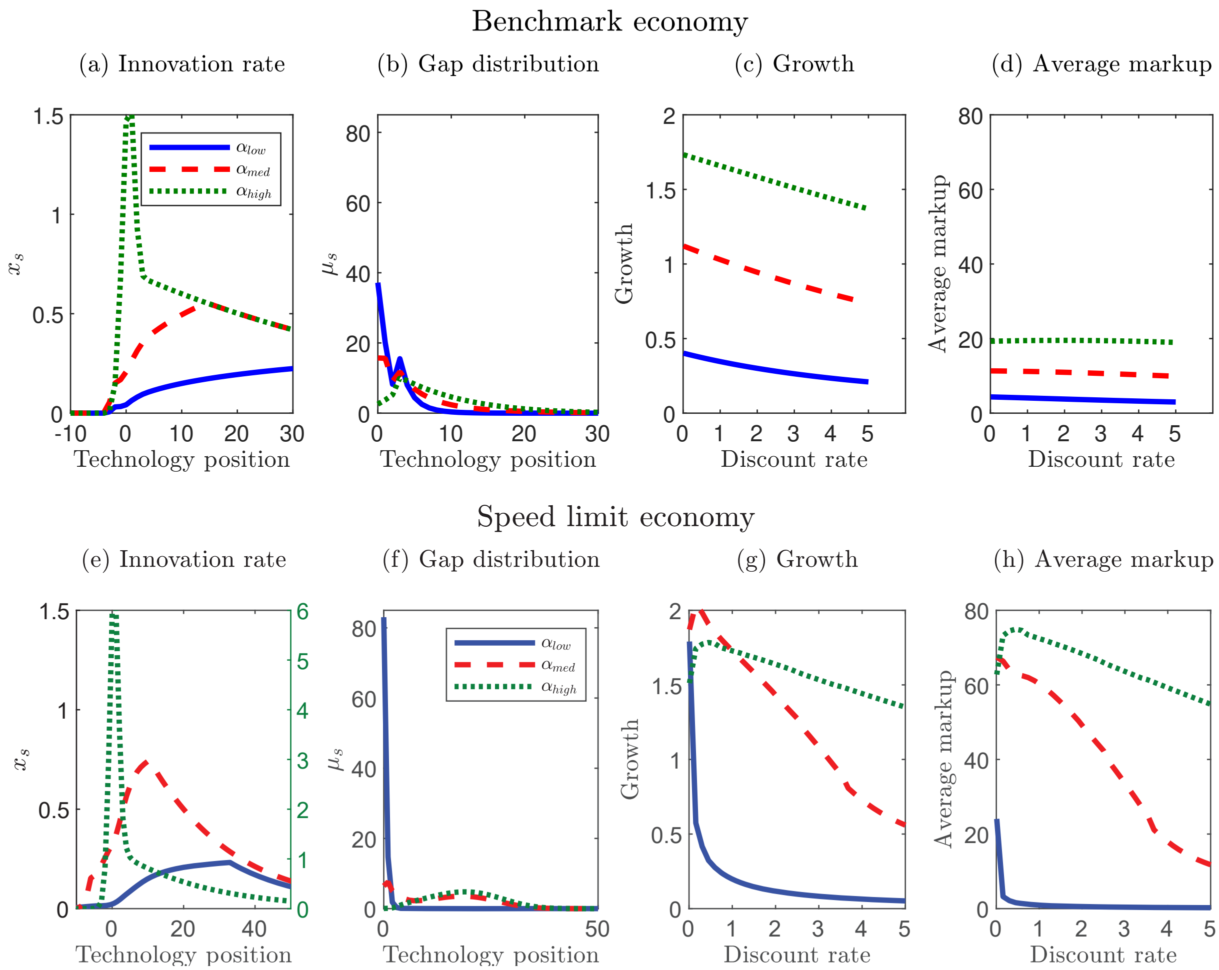

Figure 12: **Credit access.** This figure shows how credit access affects the relation of growth, competition, and the interest rate. Lower values of the limited-commitment parameter α correspond to reduced credit access.

speed limit vanishes completely: Starting from a high interest rate, a lower interest rate no longer fosters the accumulation of market power that is necessary for the growth speed-limit to obtain.

8. Policy Response to Falling Real Rate

This section performs counterfactual policy analyses to understand the implications of patent policy and R&D subsidies, and how these policies might be optimally adjusted as the interest rate varies.

8.1 Optimal uniform patent policy

We begin by asking how changing the patent expiry rate affects equilibrium outcomes. As before, we continue to restrict attention to uniform patent policies, in which the patent expiry rate does not depend on an industry's competitiveness. The black lines in Figure 13 trace out how growth and the ratio of output-to-potential vary as the patent expiry rate η changes. This frontier reflects firms' endogenous innovation strategies and the extent of market power those strategies produce. Recall, from Section 3.3, that the ratio of output-to-potential declines as market power rises, with market power both reducing the aggregate demand for labor and, when market power is heterogeneous across industries, the efficiency of its use. In each panel, the green square shows outcomes at the calibrated value of η. The dashed red lines in each panel are iso-welfare lines, showing combinations of growth and output-to-potential that deliver the same total welfare (equation 30). The slope of the iso-welfare curve is $-\rho$, capturing that growth is more important to the social planner as the discount rate falls. The figure shows this exercise in the benchmark (left panel) and sclerotic or speed-limit growth (right panel) economies.

In the benchmark economy, growth rises and output-to-potential falls as patent protection is strengthened. This result reflects a classic tradeoff for patent policy: market power incentivizes R&D but distorts production decisions (Nordhaus (1969)). With the discount rate at its calibrated value of 2 percent, the calibrated patent expiry rate is "too high": welfare is maximized by setting the expiry rate to zero. With a higher discount rate of 5 percent, the calibrated patent expiry rate is still too high, but the optimal expiry rate is now strictly positive. Thus, for a discount rate of 5 percent, outcomes under optimal uniform patent policy are shown by the tangency of the growth and output-to-potential frontier with the iso-welfare line. How optimal policy changes with the discount factor reflects changes both in the frontier of growth and output-to-potential, as well as rotation of the iso-welfare line.

In the speed-limit economy, when patent protection is very weak, the same classic

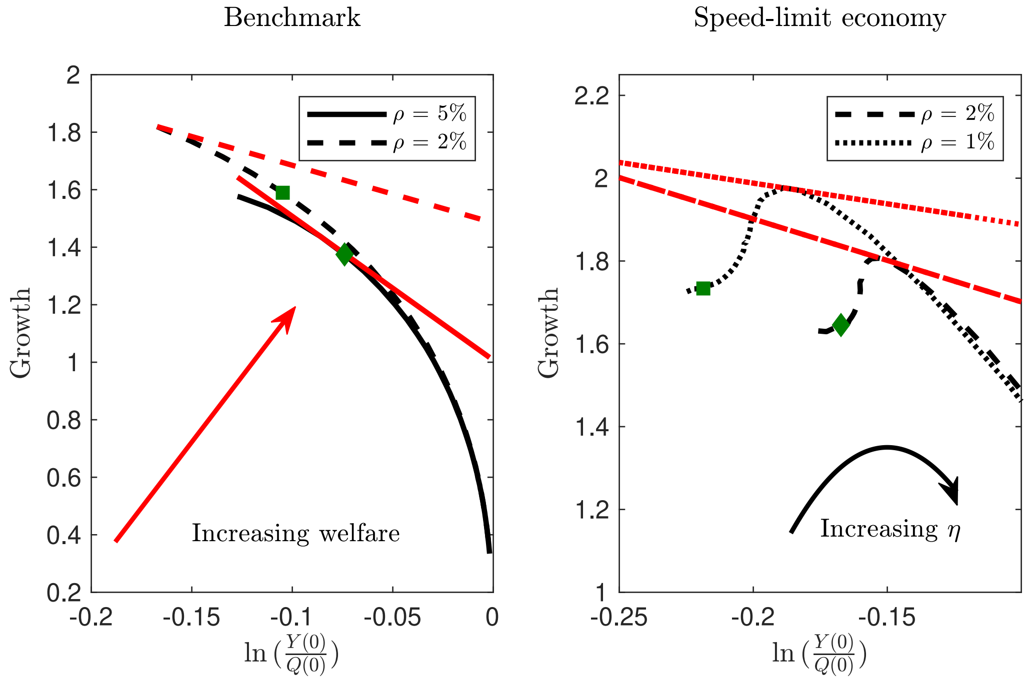

Figure 13: **Optimal uniform patent policy**. The figure traces out the frontier (in black) of growth and output-to-potential $\ln \frac{Y(0)}{Q(0)}$ as the patent expiry rate changes. Iso-welfare lines (in red) show combinations of growth and output-to-potential that provide the same level of welfare.

tradeoff arises; strengthening patent protection raises growth at the expense of higher production distortions. However, for patent expiry rates near the calibrated rate, the tradeoff disappears. Relative to the calibrated economy, optimal uniform patent policy calls for weakening patent protection, which increases growth *and* output-to-potential. This "free lunch" for the social planner generates economically important welfare gains: growth rises 17 basis points, while output-to-potential rises by 2.4 percentage points. Put differently, the social planner could maintain the same growth rate as in the calibrated economy, but by weakening patent protection, increase output-to-potential by 5 percentage points. The free-lunch gains are even larger at lower discount rates.

The top panels of Figure 14 show how the optimal uniform patent expiry rate (η) varies with the discount rate. In the benchmark economy (left panels), the optimal expiry rate is weakly decreasing as the discount rate falls. In the speed-limit economy (right panels), the opposite pattern obtains: the optimal expiry rate rises as the discount rate falls,

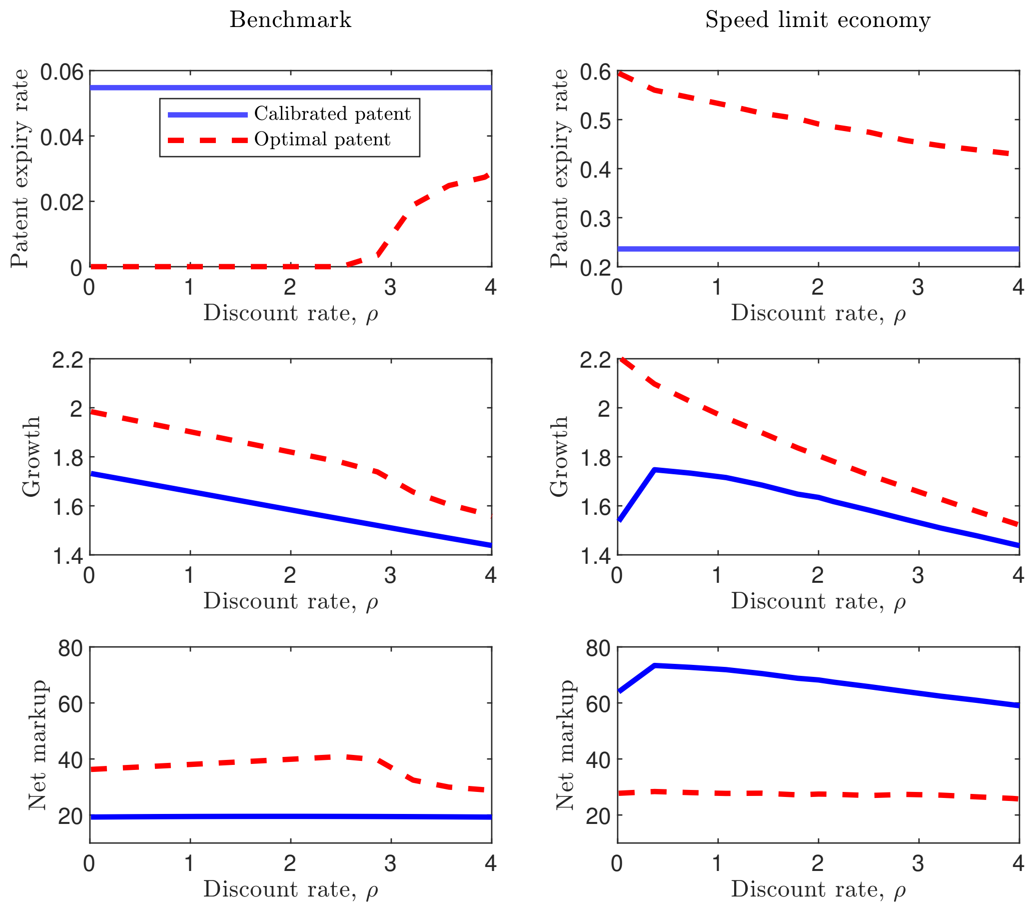

Figure 14: **Optimal uniform patent policy and its implications for growth and markups.** Solid blue lines pertain to the calibrated patent expiry rate. Dashed red lines pertain to the optimal uniform expiry rate conditional on the discount rate ρ. The discount rate is an annualized percent rate.

to counter the rise in market power. Varying the uniform patent rate optimally implies that, as the discount rate falls, there is no build-up in market power (bottom right panel). As discussed in Section 5.3, this build-up is a key part of the mechanism for the growth speed limit—the inverted-U relation between growth and the discount rate. Correspondingly, and remarkably, with optimal policy, the growth speed limit disappears (middle right panel).

Figure 15 shows the growth multiplier in the benchmark model, with the calibrated expiry rate (dashed black line) and the optimal uniform expiry rate (solid red line). Set-

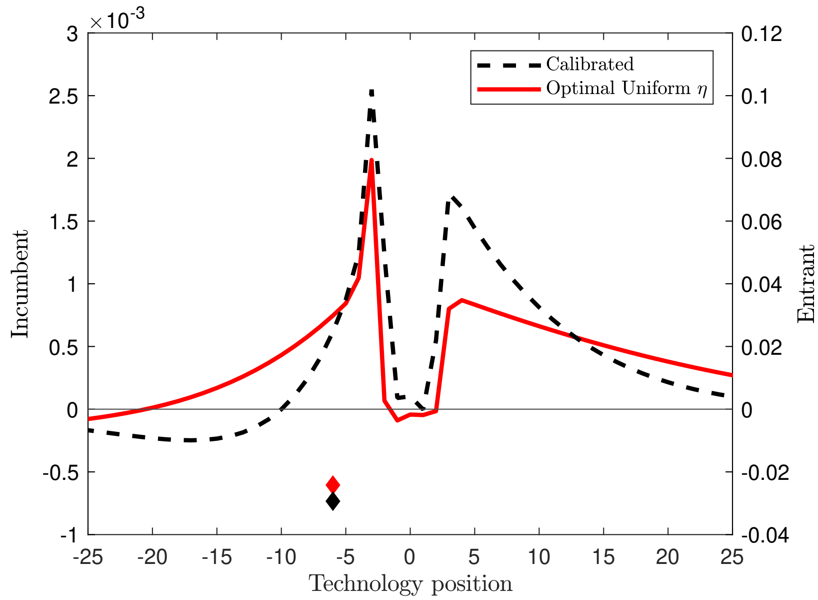

Figure 15: **The growth multiplier with the calibrated patent expiry rate and with the optimal uniform patent expiry rate.**

ting the lower, optimal patent expiry rate changes the translation of firm R&D effects into growth, in a way that depends importantly on the firm's technology position. For example, higher innovation by laggards in uncompetitive industries reduces growth in the calibrated economy but increases it in the optimal-expiry economy. This change reflects that these laggard's R&D generates a more powerful escape competition effect, and a less powerful trickle-down effect, in the optimal-expiry economy. Thus, the figure illustrates how patent policy affects the transmission of a change in the household discount rate—or of a change in R&D subsidies or patent policy—into growth.

8.2 Optimal state-dependent patent policy

State-dependent policies—which treat firms differently depending on their technology position or the competitiveness of their industry—have been shown in Schumpeterian models to offer advantages relative to uniform policies (Acemoglu et al. (2018)). Here, we study state-dependent patent policies that accords to each industry a level of patent protection that depends on the industry's competitiveness (Acemoglu and Akcigit (2012)).

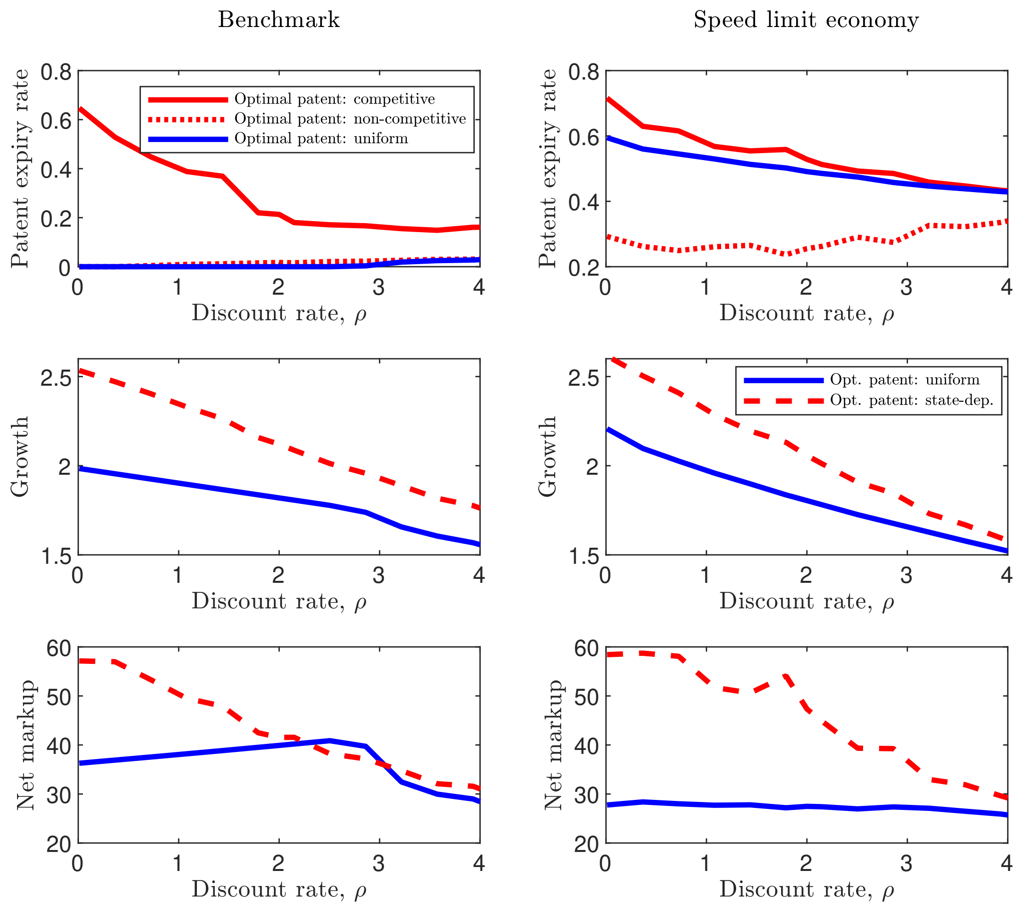

Figure 16: **Optimal state-dependent patent policy and its implications for growth and markups.** The blue lines pertain to optimal uniform patent policy. The red lines pertain to optimal state-dependent policy. In the top panel, the solid red line shows the optimal patent expiry rate for competitive industries (with small technology gaps), while the dotted red line shows the optimal expiry rate for non-competitive industries (with large technology gaps.) The discount rate is an annualized percent rate.

We illustrate industry-dependent policy by studying a "one-break" policy schedule in which industries with a technology gap above the median are given patent expiry rate η_1, while more competitive industries are given an expiry rate η_2. For a given value of the discount rate, we solve for the pair (η_1, η_2) that maximize welfare. As shown in the top panels of Figure 16, optimal state dependent policy provides much weaker patent protection in competitive industries than in non-competitive ones. Providing stronger patent protection in non-competitive industries generates a trickle-down effect, encouraging firms in competitive industries to innovate to attain a protected leadership position. Relative to optimal policy, these state-dependent policies achieve notably higher growth at the expense of higher markups. State-dependent optimal policy partly returns the build-up of market power to the sclerotic economy, but, as with optimal uniform policy, there is no growth speed limit.

8.3 Optimal Schumpeterian subsidies

R&D subsidies are a common feature of industrial policy. Here, we consider a state-dependent R&D subsidy, in which laggard firms receive a different subsidy than other firms.[41] Subsidizing laggard innovation raises interesting trade-offs. On one hand, laggard firms do not advance the technology frontier, because our calibrated parameters imply no leapfrogging. On the other hand, laggard firm innovation potentially fosters more competition, reducing production distortions and boosting growth through a composition effect. Laggard innovation also has important strategic effects, boosting leader innovation in competitive industries (escape competition) but reducing it in uncompetitive ones (trickle-down). These considerations imply that *a priori* it is unclear whether a more generous R&D subsidy for laggards—a policy we call Schumpeterian R&D subsidies—will enhance or reduce welfare.

The uniform R&D subsidy in the benchmark model (which applies to all firms) is 20 percent.[42] At a discount rate of $\rho = 2$ percent, the optimal laggard subsidy is 85 percent. Setting the laggard subsidy optimally implies a moderate increase in welfare, with an increase in growth of 10 basis points and a fiscal cost of 1.46 percent of output.[43] As the discount rate falls, the optimal laggard subsidy rises. That is, at lower discount rate, the rising optimal Schumpeterian R&D subsidy advantages laggards, whereas the falling optimal uniform patent expiry rate disadvantages them.

[41] In the model, such a policy could be implemented by identifying laggards as firms that are small with respect to revenue, profits, or employment, for example. In practice, "laggard firms" can be identified only imperfectly.

[42] These results are presented in full in the Internet Appendix.

[43] Following Atkeson and Burstein (2019), we measure fiscal cost as the change in fiscal expenditures on innovation subsidies relative to aggregate output from the initial BGP to the new BGP. We assume the policymaker has access to lump-sum transfers to focus on the economics of innovation incentives.

9. Conclusion

We developed a Schumpeterian framework to study whether lower interest rates harm innovation, long-run growth, and competition. Our framework flexibly parameterized the process of creative destruction through R&D and patent expiry. We quantified the model using indirect inference. The model matched well the cross section of markups, innovation output, and R&D and other key moments. We found that lower interest rates boost growth, with the distribution of markups little changed. These results cast doubt on recent arguments pointing to interest rates as the driver of low growth and rising markups over the past two decades.

Appendix. Model details and proofs

Value functions. For a firm in position s at time t, the (unscaled) discounted expected value of profits satisfies the Hamilton–Jacobi–Bellman equation

$$r(t)V_s(t)-\dot{V}_s(t) = \max_{x_s(t)}(1-\tau)(1-\lambda^{-s})Y(t)\mathbb{1}_{s>0}-(1-\tau_{R\&D})G(x_s(t))w(t)+x_s(t)\sum_{\hat{s}=s+1}^{\bar{s}}F_{s,\hat{s}}V_{\hat{s}}(t)-V_s(t)+$$
$$x^c_{-s}(t)\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{-s,-\hat{s}}V_{\hat{s}}(t)-V_s(t))+x_E(t)\sum_{\hat{s}=-\bar{s}}^{s-1}((\mathbb{1}_{s>0}+\frac{1}{2}\mathbb{1}_{s=0})F^E_{-s,-\hat{s}}V_{\hat{s}}(t)-V_s(t))+\eta_s\sum_{\hat{s}=0}^{s-1}(F^p_{s,\hat{s}}V_{\hat{s}}(t)-V_s(t)),$$

where $\dot{V}_s(t)$ is the derivative of $V_s(t)$ with respect to time. There is no patent expiry in tied industries: $\eta_0 = 0$. Dividing by $Y(t)$, using equations (5) and (6), and imposing that $\dot{v}_s(t) = 0$ along the BGP equilibrium, one obtains (7)–(9). A potential entrant at time t solves

$$\max_{x_E(t)} -(1-\tau_{R\&D})G_e(x_E(t))w(t) + x_E(t)\sum_{s=0}^{\bar{s}}\mu_s(t)\Bigg(\sum_{\sigma=-s+1}^{\bar{s}}F_{E-s,\sigma}V_\sigma(t)\Bigg). \tag{35}$$

Balanced growth path: Characterization. A BGP equilibrium is a vector $\Upsilon \equiv (g, \omega, \{\mu_s\}_{s\in S^+}, \{x_s\}_{s\in S}, x_E)'$ satisfying the growth equation (17) (one equation), labor market clearing (14) (one equation), incumbent and entrant first-order conditions ($2\bar{s}+1$ equations of the form (10) and one equation for the entrant, (13)), the outflow-inflow condition (19) for $s \in \{1, ..., \bar{s}\}$ ($\bar{s}$ equations), and the normalization of industry weights ($\sum_{s\in S}\mu_s = 1$, one equation).

Gap-dependent patent policy change. Patent policy is a set of gap-dependent patent expiry rates, $\{\eta_s\}_{s\in S}$. Define a gap-dependent policy change as a perturbation of this policy $\{d\eta_s\}_{s\in S}$, such that the new patent policy is $\{\eta_s + \varpi d\eta_s\}_{s\in S}$, where ϖ is a scalar and $\varpi = 0$ corresponds to the current policy.

Proof of Theorem 1. Denote $P = (\rho, \tau, \tau_{R\&D}, \varpi)$. Write the equations characterizing the balanced growth path (listed previously in this appendix, in "Balanced growth path: Characterization") as a stack $H = 0$. This stack is a set of $K = 3\bar{s}+5$ possibly non-linear equations, with the k-th equation of the form $h(\Upsilon, P; k) = 0$. The BGP equilibrium Υ satisfies $H(\Upsilon, P) = 0$. In a neighborhood of P, if H_Υ is invertible, the implicit function theorem implies

$$d\Upsilon = -H_\Upsilon^{-1} H_P dP. \tag{36}$$

H_Υ is a $K \times K$ matrix with element (k_1, k_2) corresponding to $\frac{\partial h(\Upsilon,\psi;k_1)}{\partial \Upsilon_{k_2}}$, where Υ_{k_2} is the k_2-th element of Υ.[44] H_P is a $K \times dim(P)$ matrix, where $dim(P)$ is the dimension of P. Element (k_1, k_2) of H_P is $\frac{\partial h(\Upsilon,P;k_1)}{\partial P_{k_2}}$, where P_{k_2} is the P_{k_2}-th element of P.

Consider (10), written as $G'^{-1}(\frac{\Delta v_s}{\omega}) - x_s = 0$, for $s \in S$. Using (**??**), the corresponding row of H_P is given by $\begin{bmatrix} \frac{\partial x_s}{\partial \rho} & \frac{\partial x_s}{\partial \tau} & \frac{\partial x_s}{\partial \tau_{R\&D}} & \frac{\partial x_s}{\partial \varpi} \end{bmatrix}$. Similarly, using footnote 21, the row of H_P corresponding to (13), $G_e'^{-1}(\frac{\Delta v_E}{\omega}) - x_E = 0$, is $\begin{bmatrix} \frac{\partial x_E}{\partial \rho} & \frac{\partial x_E}{\partial \tau} & \frac{\partial x_E}{\partial \tau_{R\&D}} & \frac{\partial x_E}{\partial \varpi} \end{bmatrix}$. For all other rows of H, the first three elements of the corresponding rows of H_P is a vector of zeros, because ρ, τ, and $\tau_{R\&D}$ do not enter into other equations in H. Therefore, (21) holds, with $\mathbb{M} = -H_\Upsilon^{-1} M_1$ where M_1 is a matrix that maps $\boldsymbol{x}$ into Υ.

To see that the same matrix $\mathbb{M}$ maps partial equilibrium effects of a gap-dependent patent policy change into the general equilibrium effect, consider the fourth column of H_P, which corresponds to ϖ. The row of this vector corresponding to inflow-outflow for state $s \in \{1, \ldots, \bar{s}\}$ is equal to $\sum_{\sigma \in S^+ \setminus s} \mu_\sigma \left(d\eta_\sigma F^p_{\sigma,s} \right) - (d\eta_s)\mu_s$. Stack the rows corresponding to inflow-outflow for state $s \in \{1, \ldots, \bar{s}\}$ into an $\bar{s} \times 1$ vector, $\mathbb{F}$. The remaining rows of this fourth column of H_P (corresponding to growth, labor-market clearing, and the normalization of industry weights) are equal to zero. Therefore, to first order, the effect of changing ϖ is $\boldsymbol{d\Upsilon} = \mathbb{C}_0 + \mathbb{M}\, \boldsymbol{\partial x}$, where $\mathbb{M} = -H_\Upsilon M_1$, as before, and $\mathbb{C}_0 = -H_\Upsilon M_0 \mathbb{F}$, where M_0 maps H_Υ^{-1} into the columns of H_Υ^{-1} corresponding to the inflow-outflow equations.

Proof of corollary 2. Notice that expression (36) can be written $H_\Upsilon d\Upsilon = H_P dP$. The row of this matrix equation corresponding to expression (10) provides the corollary.

Obtaining the H_Υ matrix. To obtain H_Υ, one needs values for $\{\frac{\partial v_s}{\partial \rho}\}_{s \in S}$, $\{\frac{\partial v_s}{\partial \omega}\}_{s \in S}$, $\{\frac{\partial v_s}{\partial x^c_\sigma}\}_{s \times \sigma \in S \times S}$, and $\{\frac{\partial v_s}{\partial x_E}\}_{s \in S}$. The valuation terms $\{\frac{\partial v_s}{\partial \rho}\}_{s \in S}$ are obtained as the solution to the following system of $2\bar{s}+1$ equations, which are linear in model parameters, H_Υ, and $\{v_s\}_{s \in S}$. For any

[44] In our quantitative analyses, we always find that H_Υ is invertible but we do not have a proof.

$s \in S$, taking the partial derivative of the value functions (7)-(9) with respect to ρ,

$$v_s + \rho\frac{\partial v_s}{\partial \rho} = x_s\frac{\partial \Delta v_s}{\partial \rho} + x^c_{-s}\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{-s,-\hat{s}}\frac{\partial [v_{\hat{s}} - v_s]}{\partial \rho})$$
$$+ x_E\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{E-s,-\hat{s}}\frac{\partial [(\mathbb{1}_{s>0} + \frac{1}{2}\mathbb{1}_{s=0})v_{\hat{s}} - v_s]}{\partial \rho}) + \eta_s\sum_{\hat{s}=0}^{s-1}(F^p_{s,\hat{s}}\frac{\partial [v_{\hat{s}} - v_s]}{\partial \rho}). \quad (37)$$

The wage terms $\{\frac{\partial v_s}{\partial \omega}\}_{s\in S}$ are the solution to

$$\rho\frac{\partial v_s}{\partial \omega} = x_s\frac{\partial \Delta v_s}{\partial x^c_\sigma} + x^c_{-s}\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{-s,-\hat{s}}\frac{\partial [v_{\hat{s}} - v_s]}{\partial \omega})$$
$$+ x_E\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{E-s,-\hat{s}}\frac{\partial [(\mathbb{1}_{s>0} + \frac{1}{2}\mathbb{1}_{s=0})v_{\hat{s}} - v_s]}{\partial \omega}) + \eta_s\sum_{\hat{s}=0}^{s-1}(F^p_{s,\hat{s}}\frac{\partial [v_{\hat{s}} - v_s]}{\partial \omega}) + (1-\tau_{R\&D})G(x_s). \quad (38)$$

The strategic terms $\{\frac{\partial v_s}{\partial x^c_\sigma}\}_{s\times\sigma\in S\times S}$ are the solution to

$$\rho\frac{\partial v_s}{\partial x^c_\sigma} = x_s\frac{\partial \Delta v_s}{\partial x^c_\sigma} + x^c_{-s}\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{-s,-\hat{s}}\frac{\partial [v_{\hat{s}} - v_s]}{\partial x^c_\sigma}) + x_E\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{E-s,-\hat{s}}\frac{\partial [(\mathbb{1}_{s>0} + \frac{1}{2}\mathbb{1}_{s=0})v_{\hat{s}} - v_s]}{\partial x^c_\sigma})$$
$$+ \eta_s\sum_{\hat{s}=0}^{s-1}(F^p_{s,\hat{s}}\frac{\partial [v_{\hat{s}} - v_s]}{\partial x^c_\sigma}) + \mathbb{1}_{-s=\sigma}(\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{-s,-\hat{s}}v_{\hat{s}} - v_s)). \quad (39)$$

Elasticity of intertemporal substitution. The benchmark model's assumption of log preferences can be modified to allow the EIS to differ from 1. Assume the household maximizes

$$\int_{t=0}^{\infty} e^{-\rho t}\frac{C(t)^{1-\frac{1}{\varphi}} - 1}{1-\frac{1}{\varphi}}dt. \quad (40)$$

where $\varphi > 0$ is the EIS and $\frac{1}{\varphi}$ is the coefficient of relative risk aversion. The Euler equation becomes $g = \varphi(r - \rho)$. Substituting into the firm value function, one obtains, for $s \in S$,

$$(\varphi\rho + (1-\varphi)r)v_s = \max_{x_s}(1-\tau)\mathscr{L}_s\mathbb{1}_{s>0} - (1-\tau_{R\&D})G(x_s)\omega + \Delta v_s + x^c_{-s}\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{-s,-\hat{s}}v_{\hat{s}} - v_s) +$$
$$x_E\sum_{\hat{s}=-\bar{s}}^{s-1}(F_{E-s,-\hat{s}}(\mathbb{1}_{s>0} + \frac{1}{2}\mathbb{1}_{s=0})v_{\hat{s}} - v_s) + \eta_s\sum_{\hat{s}=0}^{s-1}(F^p_{s,\hat{s}}v_{\hat{s}} - v_s). \quad (41)$$

To obtain the BGP Υ for given ρ, we search for an (r, ω) pair such that labor markets clear and the Euler equation is satisfied. Note that a BGP might not exist for sufficiently high EIS (Acemoglu and Cao (2015)).

Non-recursive value function and the duration of profits from innovation. Using the recursive Hamilton–Jacobi–Bellman equations (7)–(9), the non-recursive value function along a BGP equilibrium is:

$$v_{s(t)}(t) = \int_t^\infty e^{-\rho(z-t)}\mathbb{E}_{s(t),t}\Pi^N_{\varsigma(z)}dz, \quad (42)$$

where $\mathbb{E}_{s(t),t}$ is the expectation operator over a firm's position $\varsigma(z)$ at time $z > t$ conditional on having position $s(t)$ at time t. Net operating profits conditional on having position $\varsigma(z)$ are given by $\Pi^N_{\varsigma(z)}$.[45] The capital gain from innovation (5) can therefore be written non-recursively as

$$\Delta v_{s(t)} = \int_t^\infty e^{-\rho(z-t)}\underbrace{\left(\sum_{\hat{s}=s+1}^{\bar{s}}F_{s(t),\hat{s}}\mathbb{E}_{\hat{s},t} - \mathbb{E}_{s(t),t}\right)\Pi^N_{\varsigma(z)}}_{\text{Expected profit at time } z>t \text{ from innovation}}dz. \quad (43)$$

Using (42) and (43), we can relate the duration of profits from innovation, or $\mathscr{D}_s = -\frac{\partial \Delta v_s}{\partial \rho}\frac{1}{\Delta v_s}$, to the duration of profits, or $\mathbb{D}_s = -\frac{\partial v_s}{\partial \rho}\frac{1}{v_s}$, a widely cited metric. The relation is:

$$\mathscr{D}_s = \frac{1}{\Delta v_s}\left(\sum_{\hat{s}=s+1}^{\bar{s}}F_{s,\hat{s}}v_{\hat{s}}\mathbb{D}_{\hat{s}} - v_s\mathbb{D}_s\right). \quad (44)$$

[45]Define $\Pi^N_{\varsigma(z)} : (S \cup \varnothing) \to R$ as the net operating profit of a firm in position $\varsigma(z)$, with $\Pi^N_{\varsigma(z)} = (1-\tau)\mathscr{L}_{\varsigma(z)}\mathbb{1}_{\varsigma(z)>0} - (1-\tau_{R\&D})G(x_{\varsigma(z)})$ for $\varsigma(z) \in S$ and $\Pi^N_{\varsigma(z)} = 0$ if $\varsigma(z) = \varnothing$, with $\varnothing$ connoting that a firm has been displaced by entry prior to time z.

Perfectly elastic labor supply. Assume that the representative household maximizes

$$\int_{t=0}^{\infty} e^{-\rho t}\Big(\ln(C(t)) - L(t)\Big)dt, \tag{45}$$

where $L(t)$ is labor. As in the benchmark model, the final good is the numeraire. The household's first order conditions imply $\omega = 1$. A BGP equilibrium is, for every t, the tuple $\Upsilon \equiv (g, L, \{\mu_s\}_{s\in S^+}, \{x_s\}_{s\in S}, x_E)'$ satisfying, conditional on $\omega = 1$, the growth equation (17), incumbent and entrant first-order conditions (10) and (13)), the outflow-inflow condition (19) for $s \in \{1, ..., \bar{s}\}$, and the normalization of industry weights ($\sum_{s\in S}\mu_s = 1$), and the labor demand condition

$$L = G_E(x_E) + \sum_{s=0}^{\bar{s}} \mu_s \left[G(x_s) + G(x_{-s}) + (1 - \mathscr{L}_s)\frac{1}{\omega}\right]. \tag{46}$$

Within-industry imperfect substitution. Assume that the final good is produced using the technology

$$\ln Y(t) = \int_0^1 \ln[y_1(j)^{\frac{\kappa-1}{\kappa}} + y_2(j)^{\frac{\kappa-1}{\kappa}}]^{\frac{\kappa}{\kappa-1}} dj. \tag{47}$$

Scaled operating profits, $\pi_s^N \equiv \frac{\Pi_s^N(t)}{Y(t)}$, satisfy

$$\pi_s = \begin{cases} \frac{\nu_s^{1-\kappa}}{\kappa+\nu_s^{1-\kappa}} & \text{if } s > 0 \\ \frac{1}{\kappa\rho_{|s|}^{1-\kappa}+1} & \text{if } s > 0 \\ \frac{1}{\kappa+1} & \text{if } s = 0 \end{cases} \tag{48}$$

where ν_s is the relative price between the leader and follower in an industry with gap s. To derive (48), note that the effective demand elasticity facing firm i is $\epsilon_i \equiv -\frac{d\ln y_i}{d\ln p_i} = \kappa(1 - \delta_i) + \delta_i$, where δ_i is the market share of firm i. The markup for firm i is therefore $\epsilon_i/(\epsilon_i - 1)$. Using the linearity of each firm's production technology, one obtains that the relative price ν_s satisfies $\nu_s = \lambda^{-s}(\frac{\kappa+\nu_s^{1-\kappa}}{\kappa+\nu_s^{\kappa-1}})$. In addition, the market share of the leader firm is $\frac{\rho_s^{1-\kappa}}{1+\rho_s^{1-\kappa}}$. The labor demand of firm i is $l_i = (\delta_i - \pi_s(i))\frac{1}{\omega}$, where $s(i) \in S$ is the technology position of firm i.

References

Acemoglu, Daron and Dan Cao, "Innovation by entrants and incumbents," *Journal of Economic Theory*, 2015, *157*, 255–294.

__ and Ufuk Akcigit, "Intellectual property rights policy, competition and innovation," *Journal of the European Economic Association*, 2012, *10* (1), 1–42.

__ , __ , Harun Alp, Nicholas Bloom, and William Kerr, "Innovation, Reallocation, and Growth," *American Economic Review*, November 2018, *108* (11), 3450–91.

Aghion, Philippe and Peter Howitt, "A Model of Growth through Creative Destruction," *Econometrica*, March 1992, *60* (2), 323–351.

__ , Antonin Bergeaud, Gilbert Cette, Remy Lecat, and Hélène Maghin, "Coase Lecture - The Inverted-U Relationship Between Credit Access and Productivity Growth," *Economica*, 2019, *86* (341), 1–31.

__ , __ , Timo Boppart, Peter J Klenow, and Huiyu Li, "A Theory of Falling Growth and Rising Rents," Working Paper 26448, National Bureau of Economic Research November 2019.

__ , Christopher Harris, Peter Howitt, and John Vickers, "Competition, Imitation and Growth with Step-by-Step Innovation," *Review of Economic Studies*, 7 2001, *68* (3), 467–492.

Akcigit, Ufuk and Sina T. Ates, "What Happened to U.S. Business Dynamism?," Working Paper 25756, NBER April 2019.

__ and William R. Kerr, "Growth through Heterogeneous Innovations," *Journal of Political Economy*, 2018, *126* (4), 1374–1443.

__ , Sina T Ates, and Giammario Impullitti, "Innovation and Trade Policy in a Globalized World," Working Paper 24543, NBER April 2018.

Arrow, Kenneth, "Economic Welfare and the Allocation of Resources for Invention," in "The Rate and Direction of Inventive Activity: Economic and Social Factors," NBER, Inc, 1962, pp. 609–626.

Atkeson, Andrew and Ariel Burstein, "Aggregate Implications of Innovation Policy," *Journal of Political Economy*, 2019, *127* (6), 2625–2683.

Auclert, Adrien and Matthew Rognlie, "Inequality and Aggregate Demand," Working Paper 24280, NBER February 2018.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms*," *The Quarterly Journal of Economics*, 2 2020.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, "The Role of Entrepreneurship in US Job Creation and Economic Dynamism," *Journal of Economic Perspectives*, September 2014, *28* (3), 3–24.

Eggertsson, Gauti B, Jacob A Robbins, and Ella Getz Wold, "Kaldor and Piketty's Facts: The Rise of Monopoly Power in the United States," Working Paper 24287, NBER February 2018.

Fernald, John G., Robert E. Hall, James H. Stock, and Mark W. Watson, "The Disappointing Recovery of Output after 2009," *Brookings Papers on Economic Activity*, 2017, *48* (1 (Spring), 1–81.

Foster, Lucia, John Haltiwanger, and Chad Syverson, "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, March 2008, *98* (1), 394–425.

— , — , and C.J. Krizan, "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," in "New Developments in Productivity Analysis," NBER, 2001, pp. 303–372.

Gerschenkron, Alexander, *Economic Backwardness in Historical Perspective*, The Belknap Press of Harvard University Press, 1962.

Grossman, Gene M. and Elhanan Helpman, "Quality Ladders in the Theory of Growth," *The Review of Economic Studies*, 1991, *58* (1), 43–61.

Hall, Bronwyn, Christian Helmers, Mark Rogers, and Vania Sena, "The Choice between Formal and Informal Intellectual Property: A Review," *Journal of Economic Literature*, June 2014, *52* (2), 375–423.

Hall, Robert E, "New Evidence on the Markup of Prices over Marginal Costs and the Role of Mega-Firms in the US Economy," Working Paper 24574, NBER May 2018.

Jermann, Urban and Vincenzo Quadrini, "Macroeconomic Effects of Financial Shocks," *American Economic Review*, February 2012, *102* (1), 238–71.

Klette, Tor Jakob and Samuel Kortum, "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 2004, *112* (5), 986–1018.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," *Quarterly Journal of Economics*, 03 2017, *132* (2), 665–712.

Lentz, Rasmus and Dale T. Mortensen, "An Empirical Model of Growth Through Product Innovation," *Econometrica*, 2008, *76* (6), 1317–1373.

Liu, Ernest, Atif Mian, and Amir Sufi, "Low interest rates, market power, and productivity growth," Working Paper, Princeton University 2020.

Loecker, Jan De, Jan Eeckhout, and Gabriel Unger, "The Rise of Market Power and the Macroeconomic Implications*," *The Quarterly Journal of Economics*, 2020.

Nordhaus, W.D., *Invention Growth, and Welfare: A Theoretical Treatment of Technological Change*, MIT Press, 1969.

Peters, Michael, "Heterogeneous Markups, Growth, and Endogenous Misallocation," *Econometrica*, 2020, *88* (5), 2037–2073.