Market Power and Innovation in the Intangible Economy

Maarten De Ridder*

First version: 20 March 2019. This version: 12 July 2023

Abstract

This paper offers a unified explanation for the slowdown of productivity growth, the decline in business dynamism and the rise of market power. Using a quantitative framework, I show that the rise of intangible inputs – such as software – can explain these trends. Intangibles reduce marginal costs and raise fixed costs, which gives firms with high-intangible adoption a competitive advantage, in turn deterring other firms from entering. I structurally estimate the model on French and U.S. micro data. After initially boosting productivity, the rise of intangibles causes a decline in productivity growth, consistent with the empirical trends observed since the mid-1990s.

Keywords: Productivity, Growth, Business Dynamism, Intangible Inputs, Market Power

The decline of productivity growth has played a prominent role in recent academic and policy debates. Average productivity growth in the United States was less than 0.5% between 2005 and 2018, well below the long-term average of 1.3% (Figure 1a). A similar slowdown occurred across most of Europe, causing productivity in countries such as France and the United Kingdom to flatline (Adler et al. 2017). The slowdown followed after a decade of above-average growth, fueled by rapid improvements in information technologies (Fernald 2015). Moreover, the slowdown occurred despite an increase in productivity-enhancing investments: indeed, U.S. corporate research and development (R&D) has increased by 65% as a fraction of GDP over the last 30 years (Figure 1b). Rather than seeming to be driven by a lack of effort on the part of firms to become more productive, the slowdown can therefore be attributed to a decline in the effect of innovative investments on productivity growth (e.g. Bloom et al. 2020).

The initial surge and subsequent decline in productivity growth coincided with two other trends: the slowdown of business dynamism and the rise of markups. Signs that dynamism is weakening include the

*Department of Economics, London School of Economics, and Centre for Macroeconomics. *E-mail*: m.c.de-ridder@lse.ac.uk. *Web:* http://www.maartenderidder.com/. I am grateful to Vasco Carvalho for extensive feedback and valuable discussions. I also thank three anonymous referees and the editor, Mikhail Golosov, for their comments. I am further grateful for frequent feedback from Giancarlo Corsetti, Bill Janeway, Damjan Pfajfar and Coen Teulings, as well as for discussions with Klaus Adam, Andy Atkeson, Ariel Burstein, David Baqaee, Tiago Cavalcanti, Chiara Criscuolo, Luca Dedola, Wouter den Haan, Romain Duval, Oliver Exton, Emmanuel Farhi, Lukas Freund, Joe Hazell, Jonathan Haskel, Marie Hoerova, Hugo Hopenhayn, Peter Karadi, Joseba Martinez, Ben Moll, Dmitry Mukhin, Dimitris Papanikolaou, Michael Peters, Sophie Piton, Christopher Sandmann, Pontus Rendahl, Alex Rodnyansky, Christina Romer, David Romer, Petr Sedláček and Joshua Weiss. I thank Jean Imbs, and especially Basile Grassi, for their help with data access, and Isabelle Mejean for help with merging French datasets. This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the 'Investissements d'Avenir' program (reference: ANR-10-EQPX-1, CASD). Financial support from the Centre for Macroeconomics, Prins Bernhard Cultuurfonds and the Keynes Fund is gratefully acknowledged.

Figure 1. Trends in Productivity Growth and Research & Development

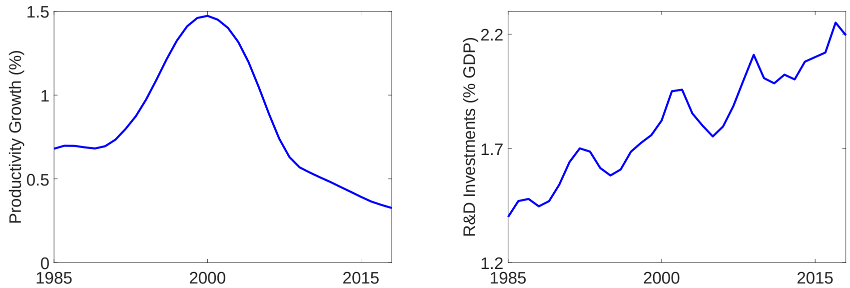

(a) Productivity Growth (b) R&D Investments

Notes: Figure 1a plots annual productivity growth from the Fernald series (Fernald 2015). The plot is smoothed using an HP filter with an annual smoothing parameter of 100. Figure 1b plots private R&D as a percentage of GDP. Data from the Bureau of Economic Analysis (2019).

decline in the rate at which workers reallocate to different firms (e.g. Decker et al. 2014), the decline in skewness of the firm-growth distribution (e.g. Decker et al. 2016) and the decline of entry rates (e.g. Pugsley and Șahin 2018). The rise of markups has recently attracted attention and has been linked to the decline of the labor share (e.g. De Loecker, Eeckhout and Unger 2020). Despite the growing body of evidence detailing these trends, there is thus far no consensus on what has caused them.

This paper claims that the trends in productivity growth, business dynamism and markups are reflections of a structural change in the way firms produce. Specifically, I show that an increase in the use of intangible inputs can drive these patterns. Intangible inputs are inputs that are used in production, but are not physically embodied. Information technology and software are prominent examples. The rise of intangible inputs has been dramatic over the last 30 years: software alone is now responsible for 18% of U.S. corporate investments, up from 3% in 1980 (Bureau of Economic Analysis 2019). These intangible inputs, typically customized for the firm that deploys them, are rarely patented, which limits positive spillovers to competitors (e.g. Bessen 2020, Aghion et al. 2022). Intangible inputs such as software are thus distinct from R&D, whose spillovers enable long-term growth (Grossman and Helpman 1991, Aghion and Howitt 1992).

Intangible inputs alter the relationship between profitability, firm-level innovation and aggregate growth because intangibles have two features: they are scalable, and firms differ in the efficiency with which they deploy them. Intangibles are scalable in the sense that they can be duplicated at close-to-zero marginal cost (e.g. Haskel and Westlake 2017, Hsieh and Rossi-Hansberg 2023). This causes the cost structure to change when firms use intangible inputs in production. Firms invest in the development and maintenance of intangible inputs but face minimal additional costs when production is scaled up. Firms that sell products that include software (e.g. the operating system of a phone, a car's drive-by-wire-system), for example, face minimal costs of reproducing software in additional units. The rise of intangibles therefore shifts costs away from the marginal towards the fixed component. As a consequence, firms take these costs as given when making decisions about pricing and production. While intangible inputs are thus non-rival in the production of a firm's goods, the limited extent of their spillovers distinguishes them from the non-rival technologies described in Romer (1990), that arise as a result of R&D.

Firms differ in the extent to which they adopt intangible inputs in order to reduce marginal costs. A 2018 European Investment Bank survey finds that over 40% of American and European manufacturing firms do not use state-of-the-art digital technologies, while less than 15% organize their entire operation around digital technologies (Veugelers, Rückert and Weiss 2020).[1] A likely driver is the fact that firms, even within narrowly defined industries, differ in the efficiency with which they are able to use intangibles. A rich literature provides evidence on this. Bloom, Sadun and Van Reenen (2012), for example, show that American-owned European establishments achieve greater productivity improvements from information technology (IT).[2] They find that IT productivity is a firm characteristic, especially because the IT productivity of establishments increases when they are *acquired* by an American firm. Schivardi and Schmitz (2019) furthermore show that inefficient management practices can explain not only the low IT adoption by Italian firms but also why the productivity gains that these firms obtain from using IT are limited.[3]

I show that a rise in the use of intangibles that is unequal across firms alters both the rate and the efficiency with which firms engage in research and development, causing the trends depicted in Figure 1. I derive these results in an endogenous growth model – in the spirit of Klette and Kortum (2004)– that is tractable yet sufficiently rich to quantitatively analyze the effect of intangibles on growth, dynamism and market power. Firms produce one or multiple goods and invest in research and development (R&D) to obtain higher-quality patents for goods produced by other firms. Successful innovation causes the innovator to become the new producer, while the incumbent ceases to produce the good. Step-wise improvements to goods through this process of creative destruction drive aggregate growth.

Firms in the model are able to reduce marginal costs by committing to the purchase of fixed-cost intangibles. As firms are heterogeneous in the efficiency with which they deploy intangibles, some firms reduce marginal costs by a greater fraction than others do. This introduces a tradeoff between *quality* – which grows as a result of R&D – and *price* – which depends on firms' use of intangibles – to the Klette and Kortum (2004) framework. In the standard model, firms that develop a higher-quality version of a good become its sole producer. Other firms have the same costs but are unable to produce the same quality, and hence cannot compete. Intangibles change this result, as high-intangible firms are able to produce at lower costs than others can – enabling them to sell at lower prices. When a firm with lower intangible-adoption develops a higher quality version of a good sold by one of these firms, the incumbent could undercut the innovator on price. Only if the quality difference is sufficiently large to offset the gap in marginal costs would the innovator become the new producer. The presence of firms with a high take-up of intangibles therefore deters other firms from entering new markets. Firms with high-intangible productivity can therefore negatively affect productivity growth.

[1] A full literature review on firm-level determinants of IT adoption is provided in Haller and Siedschlag (2011).

[2] Software and IT are used throughout this paper as examples of inputs that are both scalable and deployed at heterogeneous efficiency, as investments in these inputs have increased rapidly over the last 30 years. Other inputs that satisfy both requirements could be used to explain the trends in productivity growth, business dynamism and markups in the framework.

[3] Bloom et al. (2014) also find that structured management practices are closely related to IT adoption in American firms. Evidence also suggests that workplace organization and organizational capital affect a firm's IT productivity (e.g. Crespi, Criscuolo and Haskel 2007, Bartel, Ichniowski and Shaw 2007). Changes to organization design come at the price of high adjustment costs, which makes IT productivity a persistent firm characteristic (e.g. Bresnahan, Brynjolfsson and Hitt 2002).

It follows that the effect of the rise of intangibles on growth critically depends on how inclusive that rise is. A broad-based shift towards intangibles can raise growth because their fixed-cost nature enhances profitability and therefore incentivizes innovation – thereby stimulating growth and raising welfare. An unequal rise of intangibles, however, incentivizes innovation only for high-intangible firms, while making it harder for other firms to enter new markets. This reduces growth and welfare.

To understand the degree to which intangibles explain the macroeconomic trends, I introduce high-intangible entrants to an economy where firms initially use similar levels of intangibles. Over the transition path, high-intangible firms initially cause a boom in productivity growth. As these firms have a greater incentive to invest in R&D, they serve to "disrupt" sectors, and economic activity concentrates disproportionately around these firms. Their entry raises productivity because high-intangible firms produce at a lower cost. The increase in aggregate productivity is not matched by wages, however, because high-intangible firms set proportionally higher markups. As the economy transitions to the new balanced growth path, there is a decline in entry, as most start-ups are unable to compete with high-intangible incumbents. Low-intangible incumbents similarly have weaker incentives to innovate. This causes a gradual decline in productivity growth, which falls below the initial steady-state within 20 years after high-intangible firms first enter the market. Although overall R&D increases, it concentrates around a smaller group of incumbents. Because returns are concave, the concentration of R&D lowers its effectiveness. Combined with the fact that a fraction of innovations fail because high-intangible incumbents undercut innovators on price, this explains why growth falls while R&D increases.

The model offers three main theoretical insights. First, it shows that the relationship between aggregate R&D and aggregate growth depends on how R&D is distributed across firms. As firm-level innovation is concave in spending, concentration of R&D can negatively affect growth. As R&D concentration increases when profitability diverges across firms, heterogeneous profits come at a dynamic cost.[4] Policies such as R&D subsidies should therefore be designed with heterogeneity in firm-level incentives in mind. Second, the model introduces a tradeoff between quality and price to Schumpeterian growth models. High-intangible firms are able to sell at lower prices, which can compensate for lower quality and therefore be used to undercut innovators. Because long-term growth arises from improvements in quality, differences in intangible productivity across firms reduce the effect of R&D on growth. An increase in the variance of intangible productivity can therefore negatively affect welfare even if the variance rises due to the entry of highly productive firms. Third, the model shows that it is important to distinguish between sources of cross-sectional productivity differences that arise from productivity upon which other firms can build (denoted as quality in the model) and those that arise from technologies lacking such positive spillovers (intangibles inputs). Improvements in quality have higher public returns because future innovators stand on the shoulders of today's R&D.[5] Statically, both act as substitutes. Dynamically, substitution towards intangible inputs can carry significant welfare costs.

[4]The dynamic costs form an additional inefficiency from profit heterogeneity, on top of the static costs of heterogeneous profits and markups that arise through low production rates by high-profit firms (e.g. Edmond, Midrigan and Xu 2022, Peters 2020).

[5]This highlights an important difference between R&D, which is often referred to as intangible capital (e.g. Corrado, Hulten and Sichel 2009, McGrattan and Prescott 2014, McGrattan 2020), and intangible inputs, such as software.

I quantify the model using two structural estimations, one for the U.S. and one for France. The French estimation relies on administrative data for the universe of firms while its U.S. counterpart relies on data for listed firms. While evidence on the macroeconomic trends is stronger for the U.S., I show that the trends are largely visible for France as well. Using a new measure of fixed costs, I show that the share of fixed over total costs gradually rose from 13 to 23% in the U.S. between 1979 and 2015, and from 9 to 14% in France between 1994 and 2016. I also use the micro data to confirm model-implied conditional correlations between intangibles, fixed costs, markups, and innovation.

The quantitative model explains a significant part of the slowdown of productivity growth, the decline in business dynamism and the rise of markups. The model predicts a slowdown in steady-state growth of 0.3 percentage points in the U.S., after an initial boom in growth of eight years. For France, the model's predictions are more modest, with a 0.1 percentage point slowdown. Markups increase by 14.6 and 5.6 percentage points in the respective calibrations. Entry rates fall by 4.6 and 1.1 percentage points, respectively. If markups are held constant, the model predicts a greater decline in productivity growth and business dynamism. The rise of markups stimulates innovative investments by high-intangible firms, and therefore mitigates the decline in productivity growth in the baseline model.[6]

Related literature This paper contributes a new mechanism to the literature that links trends in productivity growth, business dynamism and market power. Closest to this paper is Aghion et al. (2023), who also point to technology as the driver of the trends. In a model with creative destruction, they analyse the effect of an increase in span of control, which enables firms with higher (ex-ante) production efficiency to expand. Their short-term predictions overlap with mine: the efficient firms expand and cause a burst in growth, while aggregate markups rise because economic activity reallocates to profitable firms. In the long run, Aghion et al. predict that efficient firms face increasing competition from other efficient firms. This diminishes their incentives to innovate, so that growth and investments in R&D decline.[7] In my model, high-intangible firms invest persistently more in R&D, as they remain more profitable than their competitors. The additional investments have low returns because innovation is concave in a firm's R&D. As high-intangible firms can undercut others, they lower both the level and the effectiveness of R&D by other firms. Intangible inputs therefore help us to understand why the slowdown of productivity growth occurred despite a continued rise in corporate R&D.

Alternative explanations of the macroeconomic trends focus on the long-term decline in productivity growth, rather than the initial burst and subsequent slowdown. Akcigit and Ates (2022, 2021) hypothesize that imitation rates between leaders and followers have declined, and provide evidence that intellectual property rights are increasingly used anti-competitively. Olmstead-Rumsey (2022) documents that the innovation efficiency of laggard firms has fallen over time. This causes a rise in markups and a decline in aggregate R&D in a quantitative growth model. I contribute to their findings by offering an endogenous driver of the decline in the innovation efficiency of laggards: as high-intangible firms

[6]The analysis is therefore robust to concerns about the firm-level measurement of markup, on which Syverson (2019), Basu (2019) and Bond et al. (2021) provide careful discussions.

[7]Profitability and the incentives to innovate fall over time because a greater fraction of producers (firms with the highest quality patent for a product) face second-best competitors that also have a high production efficiency. Under Bertrand competition, this reduces the limit-price markup for the producing firm, so that the incentive to innovate falls.

become dominant, low-intangible (laggard) firms are more likely to be undercut when they try to enter a new market, thereby lowering the returns to their research activities.

I also view my paper as complementary to the growing body of work that explains the long-term decline in productivity growth through demographics. Peters and Walsh (2021) relate the decline of entry to the fall in labor force growth, which is consistent with evidence in Hopenhayn, Neira and Singhania (2022) and Karahan, Pugsley and Şahin (2019). The lack of entry stimulates expansion by incumbents, which raises firm concentration and markups, and slows down productivity growth.[8] Liu, Mian and Sufi (2022) relate productivity and business dynamism to low interest rates, which may itself be caused by aging (e.g. Eggertsson, Mehrotra and Robbins 2019). Low interest rates increase investment most strongly for the market leader which, through strategic interaction, dissuades followers from investing in R&D. This in turn reduces optimal R&D by leaders, and therefore slows productivity growth.[9] I offer a complementary mechanism that combines these paper's predictions of a long-term slowdown in growth with an increase in aggregate R&D.[10]

This paper also provides a theory of the trends in productivity growth, business dynamism and market power that is consistent with their micro properties. As emphasized by Van Reenen (2018), there has been a substantial divergence in profitability and productivity across firms over time. Andrews et al. (2016) show that productivity growth of the most productive firms has not declined. Decker et al. (2020) find an increase in productivity dispersion within the U.S.[11] The rise in markups in De Loecker, Eeckhout and Unger (2020) is also strongest in the highest deciles.[12] I propose that intangible inputs are at the root of the growing differences across firms, and show that the negative externality that high-intangible firms impose on others can drive aggregate trends in productivity growth, business dynamism and markups.

This paper is additionally consistent with Baqaee and Farhi (2020)'s finding that aggregate markups have increased because output has reallocated towards high-markup firms. Their finding, in turn, aligns with the finding of Autor et al. (2020) and Kehrig and Vincent (2021) that the decline in the labor share is driven by reallocation. Most other theories of rising profits and falling growth predict, instead, a within-firm rise of markups.[13] I predict that high-intangible, high-markup firms endogenously expand through R&D. Even though firm-level markups are constant, aggregate markups rise over time.

Other papers also relate the rise of intangible inputs to firm concentration. Hsieh and Rossi-Hansberg (2023) suggest that intangibles explain the rise of concentration in services across geographic markets,

[8] Peters and Walsh (2021) note that, with an alternative calibration, their model can predict both a rise in R&D and a decline in growth after a decline in population growth. Salgado (2020) notes that there has also been a decline in the fraction of college graduates that become entrepreneurs. Bornstein (2018) relates the rise of markups to aging and price sensitivity.

[9] Liu, Mian and Sufi (2022) predict that a shock to interest rates initially causes a burst of productivity growth, but a slowdown follows immediately and productivity growth falls below the initial steady-state level within a quarter.

[10] Cavenaile, Celik and Tian (2020) extend the Schumpeterian growth model with oligopolistic competition. Brynjolfsson, Rock and Syverson (2021) claim that adopting AI requires unmeasured investments, causing productivity to initially decline but eventually rise. Rachel (2021) notes that leisure-enhancing technologies can exacerbate the understatement of utility in observed TFP.

[11] Kehrig and Vincent (2019) note that an increase in productivity dispersion at the establishment level may reflect an improvement in factor allocation and a reduction of internal credit market frictions.

[12] This result has been confirmed for several countries (Díez, Fan and Villegas-Sánchez 2021, Calligaris et al. 2018). Altomonte et al. (2021) remark that credit frictions can cause dispersion in intangibles and markups across firms.

[13] Aghion et al. (2023) is a notable exception. Theories of rising profits through reallocation outside of the endogenous growth literature include Helpman and Niswonger (2022), Hubmer and Restrepo (2021) and De Loecker, Eeckhout and Mongey (2021).

as software can be deployed across markets after paying a fixed cost.[14] Korinek and Ng (2019) and Weiss (2020) relate fixed-cost intangibles to the rise of firm concentration and markups. Lashkari, Bauer and Boussard (2022) share these predictions and provide evidence on the non-homothetic role of IT in production.[15] Mariscal (2018) adds that IT complements the efficiency of managers, and may explain wage polarization. I view my paper as complementary, as I show that fixed-cost intangibles can also explain the trends in productivity growth: its initial burst and subsequent slowdown, as well as the simultaneous rise in R&D.

Related to this, the paper highlights that intangibles such as software are conceptually different from product-enhancing R&D in terms of their contribution to long-term growth. This relates to Levin et al.'s (1987) idea that process improvements carry limited spillovers and are often not patented, a notion that has since been frequently validated (see Aghion et al. 2022 for a review). While intangible inputs and R&D are sometimes studied jointly (e.g. Corrado, Hulten and Sichel 2009), their economic properties are thus distinct.

My theoretical predictions are in line with empirical work that relates productivity growth, business dynamism and market power to intangibles. Crouzet and Eberly (2019) show that intangibles cause an increase in market power and productivity for leading U.S. public firms. McKinsey (2018) and Ayyagari, Demirguc-Kunt and Maksimovic (2018) show that firms with high profitability and growth invest more in software and R&D. Bessen and Righi (2019) find that productivity of U.S. firms increases persistently after an increase in the stock of their IT staff. Farhi and Gourio (2018) show that unmeasured intangibles can explain the rising wedge between the measured marginal product of capital and risk-free rates. Bajgar, Criscuolo and Timmis (2019) find that sectors with high intangible investments experienced a greater increase in concentration. Bessen (2020) finds a positive sector-level relationship between concentration and the use of IT systems, and stresses that the scalability of intangibles is advantageous to firms that are already large. Calligaris et al. (2018) find a positive correlation between the use of digital technologies and the rise of markups. Bijnens and Konings (2018) find that the decline in Belgian business dynamism is strongest in high-IT industries.[16]

The theoretical framework builds on Schumpeterian growth models of creative destruction in the tradition of Aghion and Howitt (1992) and Grossman and Helpman (1991). It is part of the strand of models where firms produce multiple products (Klette and Kortum 2004). The Klette and Kortum framework is attractive because it is analytically tractable, yet able to replicate many empirical features of firm dynamics (Lentz and Mortensen 2008). The framework was recently used to study the reallocation of innovative activity (Acemoglu et al. 2018), to discern the effect of innovation policy (Atkeson and Burstein 2019) and to assess sources of innovation (Akcigit and Kerr 2018, Garcia-Macia, Hsieh and Klenow 2019). It has also been used to analyze misallocation with endogenously heterogeneous markups (Peters 2020).

[14] Babina et al. (2020) find evidence that artificial intelligence has a similar effect on concentration.

[15] Martinez (2019) and Martinez (2021) relate automation to the labor share.

[16] The paper also relates to the literature on rising profits (e.g. Gutiérrez and Philippon 2019, Barkai 2020, Karabarbounis and Neiman 2019) and how measures of profits may relate to intangibles (e.g. Koh, Santaeulàlia-Llopis and Zheng 2020).

Outline The paper proceeds as follows. Section I presents the growth model and discusses the mechanism. Section II verifies the model's main predictions empirically. I structurally estimate the model in Section III, and discuss results in Section IV. Section V presents extensions, and Section VI concludes.

I. Intangibles, Firm Dynamics and Growth

This section describes the endogenous growth framework with creative destruction and intangible inputs that I use to explain the trends in productivity growth, business dynamism and market power.

A. Preferences and Market Structure

There is a continuum of identical households with unit mass that choose the path of consumption to maximize the following utility function:

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

where C_t is consumption and ρ is the discount factor. Time is continuous and indexed by t, which is suppressed when convenient. The household is endowed with a single unit of labor, which it supplies inelastically. The consumption good is composed of a continuum of intermediate goods, indexed by j. Each good can be produced by the set of firms I_j that own the production technology, a patent, to produce good j at a level of quality q_{ij}. Quality determines the value that each unit of a good produced by a firm $i \in I_j$ contributes to aggregate consumption. The intermediate goods are competitively aggregated with the following Cobb-Douglas technology:

$$Y = \exp\left(\int_0^1 \ln\left(\sum_{i \in I_j} q_{ij} y_{ij}\right) dj\right),$$

where Y denotes aggregate output, and y_{ij} is the amount of good j that is produced by firm i. As all output is consumed, $Y = C$. Firms that own the patent to produce good j compete à la Bertrand. While multiple firms own such patents, consumers only demand j from the firm that offers the highest combination of output and quality at a given expenditure. In other words, goods are produced by the firm that offers the lowest quality-adjusted price p_{ij}/q_{ij}.

B. Firms and Intangibles

There is a continuum of firms, indexed by i. In the spirit of Klette and Kortum (2004), firms potentially produce more than one good, as they can produce any good for which they own a patent.

Firms produce each of their goods using two inputs, a tangible and an intangible input. The intangible input is an input that allows firms to reduce a good's marginal cost by some desired fraction. To preserve tractability, the only tangible input is production labor l_{ij}, so that intangibles allow firms to cut the amount of production labor required to produce an additional unit of output. Denoting by $s_{ij} \in (0,1]$ the fraction of marginal costs that a firm incurs, the production function reads

$$y_{ij} = l_{ij}/s_{ij}, \tag{2}$$

so that the marginal cost for firm i of producing j equals $mc_{ij} = s_{ij}w$, where w is the wage rate.

To lower s_{ij} in order to reduce marginal costs, firms must raise spending on intangible inputs. Intangibles differ from production labor in two ways. First, firms differ in the efficiency with which they deploy intangibles. A firm-specific cost parameter, ϕ_i, determines by how much a firm's marginal costs fall for a given expenditure on intangibles. This causes some firms to produce more efficiently than others. Firms draw their ϕ_i from a known discrete distribution $G(\phi)$ at birth and benefit from their intangible efficiency on each good that they produce. Second, firms commit to their spending on intangibles before they observe the marginal costs of competitors and before they set prices. Their intangibles are sunk when firms make pricing and production decisions; firms incur these costs regardless of how much they produce, which is why intangibles represent a fixed cost. Note that firms pay the fixed costs at each t. This keeps the model's tradeoffs extremely tractable, and limits the difference between tangible and intangible inputs to the minimum needed for the results.[17] Additionally, recurrent costs are in line with the high depreciation rate of software, which Li and Hall (2020) estimate at 30 to 40%. Firms must therefore constantly invest in order to maintain constant levels of software.[18]

Total spending on intangible inputs to achieve a marginal cost of $s_{ij}w$ is given by the function

$$f(s_{ij},\phi_i) = w\phi_i \left(s_{ij}^{-\theta} - 1 \right), \tag{3}$$

where $\theta > 0$. The function declines convexly in $s_{ij} \in (0,1]$, such that firms must raise intangible spending to achieve lower marginal costs for the product. The function rises in $\phi_i > 0$, such that firms that are efficient at using intangibles spend less to achieve a certain marginal cost. The function furthermore satisfies $f(1,\phi_i) = 0$, so that firms do not pay for intangibles if they do not reduce marginal costs, and $\lim_{s_{ij} \to 0} f(s_{ij},\phi_i) = \infty$, which assures that all firms have positive marginal costs in equilibrium.

The timing of production decisions is as follows. Firms first observe the intangible costs and quality levels of their competitors for the products that they can produce. For each product, they then separately choose whether to produce at the baseline marginal costs w, or to retain – by paying a vanishingly small sunk cost – the option to reduce marginal costs by spending on intangibles.[19] Firms then choose their marginal costs and pay the associated fixed costs on intangibles. Finally, firms observe their competitor's marginal costs, set prices, and produce the quantity of each product that is demanded from them.

[17] Note that firms in the model also accumulate intangible capital in the spirit of Corrado, Hulten and Sichel (2009): they invest in research and development to expand the set of goods that they produce, which persistently affects both firm size and GDP.

[18] There has furthermore been an increase in the share of enterprise software that is sold *as a service* (SaaS), where firms pay periodic fees instead of one-off fees for perpetual use. For example, 35% of Microsoft's enterprise sales in Q2 of 2019 came from SaaS, at an annual growth rate of 48%.

[19] The small cost $\epsilon \to 0$ that enables firms to choose positive intangibles represents, e.g., the cost of setting up an IT department. A similar assumption simplifies Nash equilibrium pricing in Akcigit and Kerr (2018) and Acemoglu et al. (2018).

C. Static Equilibrium

Before proceeding to how firms obtain new patents to produce additional goods, consider the static equilibrium where the set of firms with a patent to produce each good is taken as given. The equilibrium's main novelty is the identity of a good's producer, which will be determined not only by the quality of the firms' patents, but also by the efficiency with which they use intangibles inputs.

In the baseline Klette and Kortum (2004) model, the firm with the highest quality patent is the sole producer. This is because firms have equal marginal costs, so that the highest quality producer can offer the lowest price per quality unit. In the present model, marginal costs differ across firms through their choice of intangibles, the fixed costs of which are determined by their intangible cost parameter ψ_i. This means that firms can sell at different prices, which can compensate for lower levels of quality. As I show formally below, the equilibrium producer is the firm which – after adjusting for quality differences – has the lowest *choke price*. This choke price is the price at which a firm breaks even, were it to optimize intangibles as the sole producer selling at that price. Firms with a relatively low ψ_i have a lower choke price, which means that they can reduce marginal costs by a greater fraction than other firms can while still covering their fixed costs. A firm with a lower quality patent may thus still be able to offer the best combination of prices and quality, and therefore to produce in equilibrium.

To derive this subgame-perfect Nash equilibrium, note that an equilibrium satisfies two necessary conditions. First, there is only one firm $i \in I_j$ that actively produces good j in any subgame perfect equilibrium. If there were more than one firm, ruinous Bertrand competition would imply that at least one firm ends up setting prices equal to marginal costs in the pricing stage. As intangible inputs are sunk costs at that stage, that firm would produce at a loss. Second, the unique active firm in I_j is the firm that can propose the lowest profitable yet entry-proof quality-adjusted price. Formally, define firm i's choke price as the lowest price at which it does not incur a loss if it is the good's sole producer:

$$p_i^c = \inf\left\{ p \geq 0 : \max_{s_{ij} \in (0,1]} (p - w s_{ij}) Y p^{-1} - f(s_{ij}, \psi_i) \geq 0 \right\},$$

which uses that demand for the producer is $y_{ij} = Y p_{ij}^{-1}$. Because demand does not depend on quality once a good's producer is identified, a firm's choke price is independent of product characteristics and solely determined by the firm's intangible cost parameter – in which the choke price monotonically increases (see Appendix A.1). We can therefore write $p_i^c = p^c(\psi_i)$. By construction, if for two competing firms i and $\tilde{i}$, we have that i has a lower quality-adjusted choke price, then firm i can price $\tilde{i}$ out of the market at any price at which $\tilde{i}$ is willing to be active. Hence firm $\tilde{i}$ avoids incurring sunk costs on intangibles, thus never setting $s_{\tilde{i}j} < 1$. It follows that the single active firm, henceforth firm i, satisfies

$$\frac{p^c(\psi_i)}{q_{ij}} = \min_{h \in I_j} \frac{p^c(\psi_h)}{q_{hj}}. \tag{4}$$

Note that, while the level of quality does not affect how much of good j is demanded from the producer for a given price, quality is key to determine the identity of the equilibrium producer.

It remains for us to characterize the equilibrium price and marginal costs of the producer. Because quality units by the potential producers of a good are perfect substitutes, the producer's price is constrained by the marginal costs of its nearest competitor. Anticipating that they will be undercut, the producer's competitors do not invest in intangibles. This means that, in line with the standard model, the equilibrium price is equal to the wage – adjusted for the quality of the closest firm's patent.

Formally, the producer charges the following price in the final stage of the interaction:

$$p^*_{ij} = w \times \frac{q_{ij}}{\max\limits_{\tilde{i} \in \tilde{I}_j} q_{\tilde{i}j}}, \tag{5}$$

where $\tilde{I}_j = I_j \setminus \{i\}$ are firm i's competitors in the production of good j. Anticipating that firms in $\tilde{I}_j$ do not purchase fixed-cost intangibles, the minimum price at which they are willing to sell is w. The firm in the most advantageous position to compete with firm i is therefore the firm with the highest level of quality, irrespective of that firm's choke price. The price (5) therefore deters any competitor from undercutting i – which means that the producer engages in limit pricing. Note that the ratio in (5) always exceeds one if firms have equal intangible efficiencies, as the firm with the highest quality would inherently have the lowest quality-adjusted choke price (as in Klette and Kortum 2004). If firms differ in their intangible efficiency, the quality of producer i may be lower than the highest quality among its competitors.

As demand for the producer's output depends on the price it sets, the firm internalizes (5) when it backwardly induces optimal fixed costs in the intangibles-setting stage. Minimizing the sum of variable costs $ws_{ij}y_{ij}$ and fixed costs (3) gives the following product-specific first-order condition:

$$s^*_{ij} = \min\left\{\left(p^*_{ij}\frac{1}{Y}\theta\phi_i\right)^{\frac{1}{\theta-1}}, 1\right\}. \tag{6}$$

Firms with lower intangible cost parameters ϕ_i thus incur a smaller fraction of their marginal costs although their exact marginal costs differ per product, as prices p^*_{ij} depend on the firm's relative quality.

Finally, the initial stage of the strategic interaction – in which firms pay a small sunk cost to enable the use of intangibles – ensures that the equilibrium is unique. While it is necessary for a Nash equilibrium to satisfy that (4) identifies the producer, which sets prices and intangibles in line with (5) and (6), the initial stage ensures that other firms do not use intangibles. In the absence of the small sunk cost, competitors $\tilde{i} \in \tilde{I}_j$ may find it profitable to set $s_{\tilde{i}j} < 1$ in response to (6). The initial stage rules out this indeterminacy, as competitors anticipate that they will not produce and therefore avoid sunk costs.

Combined with the first-order conditions, it follows that optimal intangibles, markups and profits can be written in terms of the producer's intangible efficiency and the producer's quality relative to its highest-quality competitor. Labeling the latter λ_{ij}, the optimal price reads $p^*_{ij} = w\lambda_{ij}$, which in turn yields that (6) can be written as

$$s^*_{ij} = \min\left[\left(\lambda_{ij}\frac{w}{Y}\theta\phi_i\right)^{\frac{1}{\theta-1}}, 1\right]. \tag{7}$$

The markup μ_{ij} is then found by dividing the price p^*_{ij} by firm i's optimal marginal cost ws^*_{ij}:

$$\mu_{ij} = \lambda_{ij}/s^*_{ij}, \tag{8}$$

which means that markups increase in λ_{ij}, as well as the producer's use of intangibles. For comparison, in the baseline Klette and Kortum (2004) model without intangibles, the markup is solely determined by λ_{ij}. That would be the case here if $s^*_{ij} = 1$. Note that profits π_{ij} do not rise proportionally to markups when firms use more intangibles, as firms also face higher fixed costs. As a result, profits are given by

$$\pi_{ij} = \left(1 - \mu_{ij}^{-1}\right) Y - f(s^*_{ij}, \phi_i), \tag{9}$$

where the first term represents variable operating profits, while the second term represents fixed costs. As the firm's s^*_{ij} is pinned down by the combination of ϕ_i and λ_{ij}, profits can be written as $\pi(\phi_i, \lambda_{ij})$.

D. Innovation

Research and Development

Firms expand their portfolio of patents by investing in research and development (R&D). When investing, firms choose the Poisson flow rate $x_i \geq 0$ with which a new patent is added to their portfolio. In exchange for achieving x_i, firms employ rd^x researchers along

$$rd^x(x_i, n_i) = \eta^x x_i^{\psi^x} n_i^{-\sigma}, \tag{10}$$

where $\psi^x > 1$ and $0 \leq \sigma \leq \psi^x - 1$. The number of researchers that the firm employs is convex in the rate of innovation and declines in the number of goods that the firm produces, n_i. The former implies that the marginal return to R&D is diminishing within each time t. The latter is an assumption from Klette and Kortum (2004), and reflects the idea that large firms have more in-house knowledge or organizational capital than small firms do. Practically, the presence of $n_i^{-\sigma}$ governs the relationship between firm size and firm growth. For $\sigma = \psi^x - 1$, the model satisfies Gibrat's law of constant firm growth in size, while for $\sigma = 0$ firm growth declines rapidly in size. Following Akcigit and Kerr (2018), I allow for an intermediate case $\sigma \in [0, \psi^x - 1]$ so that the relationship between size and growth matches the data.

A firm that innovates successfully becomes the owner of a state-of-the-art patent for a random good j. Innovation is not directed, in the sense that firms are equally likely to innovate on all products. As in Aghion and Howitt (1992), the state-of-the-art patent allows firm i to produce its new good at a quality level that is a multiple $\lambda_{ij} > 1$ of the level of the current producer of the good:

$$q_{ij} = q_{\tilde{i}j} \lambda_{ij}, \tag{11}$$

where $\tilde{i}$ denotes the incumbent of good j. The innovation step λ_{ij} is a continuous random variable with counter-cumulative distribution function $H(\lambda)$. The common notation with λ_{ij} in Section C is deliberate, as the next paragraph explains that the innovation step will also equal the difference between the quality of the equilibrium producer and the highest quality among the producer's competitors.

Innovation in the model is different from the usual Klette and Kortum (2004) setup because the innovator of a certain good will not necessarily become its new producer. Section C shows that, here, the innovator only becomes the new producer of the good if its quality improvement λ_{ij} is sufficiently large for the innovator to become the firm with the lowest quality-adjusted choke price. It follows that when the innovator is firm i and the incumbent is $\tilde{i}$, the innovation is successful if

$$\text{¸}_{ij} \, \text{È} \, \frac{p^c(\text{Á}_i)}{p^c(\text{Á}_{\tilde{i}})}. \tag{12}$$

When this condition is satisfied, the innovator receives its patent and becomes the firm with the lowest quality-adjusted choke price. This always happens if the innovator has an equal or lower choke price than the incumbent does, as it can offer the good at the same price but at a superior quality. If the incumbent has a lower choke price (implying that it uses intangibles more efficiently), the innovator must have a sufficiently large innovation step in order to take over. In case the innovator's quality improvement is also insufficiently large, the innovation fails; product j is not added to the innovator's portfolio J_i and the innovation is lost. As the innovator never produces good j, future innovators are unable to learn from the lost innovation and only improve j's quality over the level at which the incumbent produces.

It follows from (12) that a good's producer is the firm that most recently implemented an innovation. As innovators that fail to take over production lose their patent, the difference between the quality of the producing firm and that of the firm that is closest on the quality ladder is ¸_{ij}. Hence the step size ¸_{ij} is also the quality difference that determines the producer's optimal intangibles in (7) and its markup in (8).

Innovation and Intangibles

It is useful to highlight the difference between quality and price in the model. In most models of growth through creative destruction, the two are isomorphic. Prices reflect the ability of firms to produce at low marginal costs (that is, with high productivity). It may seem that this is equivalent to quality, in the sense that a firm can increase its effective output $q_{ij} y_{ij}$ using the same quantity of tangible inputs by either selling at higher quality or by using a greater amount of intangibles.

The difference between the two lies in their contribution to long-term growth. Innovation raises the quality with which good j can be produced. If an innovating firm successfully takes over production, this offers a positive externality: all future innovations on j are improvements over q_{ij}; the innovation by firm i allows good j to be produced at a permanently higher level of quality. This makes quality improvements the source of long-term economic growth. Intangibles do not come with a similar externality. They improve production efficiency only for the current producer. Intuitively, the fact that the incumbent is efficient at using software applications to reduce marginal costs does not benefit an innovating firm when it takes over production.

E. Entry and Exit

There is a unit mass of entrepreneurs that invest in R&D to obtain patents to produce goods that are currently owned by incumbents. The R&D cost function is analogous to that of incumbents:

$$rd^e(e) \, \text{Æ} \, \text{´}^e e^{\text{Ã}^e}, \tag{13}$$

where $rd^e(e)$ denotes the number of researchers employed by potential entrants to achieve entry rate e, and where $\eta^e > 0$, $\psi^e > 1$. Entrants that draw an innovation improve the quality of a random good that is currently produced by an incumbent. In similar spirit to models where firms draw idiosyncratic productivities at birth (e.g. Hopenhayn 1992, Melitz 2003), entrants then draw their $\phi_e \in \Phi$ from the known distribution $G(\phi)$, and learn about their incumbent's intangible costs. The entrant becomes the new producer if it has drawn a sufficiently large step to overcome any difference in choke prices.

A firm exits the economy if it does not produce any good in its patent portfolio. This happens when entrants or other incumbents develop higher-quality versions of the sole good that a firm produces.

F. Creative Destruction

Firms cease to produce a good if a different incumbent or an entrant successfully innovates on that product. The rate at which this happens is the rate of creative destruction, $\tau(\phi_i)$. The rate of creative destruction is endogenous, as it is determined by the respective efforts that incumbents and entrants put into innovation. It is a function of the firm's intangible cost parameter ϕ_i, because a firm with relatively low intangible costs is more likely to be able to undercut an innovative challenger on price. The rate of creative destruction for a firm with intangible cost ϕ_i is given by

$$\tau(\phi_i) = \sum_{\phi_h \in \Phi} \text{Prob}\left(\lambda_{hj} \geq \frac{p^c(\phi_h)}{p^c(\phi_i)}\right) \cdot \left[\sum_{n=1} M_n(\phi_h) x_n(\phi_h) + eG(\phi_h)\right], \tag{14}$$

where $x_n(\phi_h)$ and $M_n(\phi_h)$ respectively denote optimal innovation rates and the measure of firms with intangible cost ϕ_h that produce n products. The outer-summation reflects that an incumbent with intangible cost ϕ_i faces innovative competitors from each intangible-cost level $\phi_h \in \Phi$. Within the summation there are two terms: the probability that an innovation by a firm with cost ϕ_h is successful, multiplied by innovation efforts by firms with that level of intangible costs. Under cumulative density function $1 - H(\lambda)$, the probability that (12) is satisfied when i is the incumbent and h innovates equals

$$\text{Prob}\left(\lambda_{hj} \geq \frac{p^c(\phi_h)}{p^c(\phi_i)}\right) = H\left(\frac{p^c(\phi_h)}{p^c(\phi_i)}\right). \tag{15}$$

This probability is strictly lower when the incumbent is a low-ϕ firm, as these have a lower choke price. The term for innovation effort in (14) contains two parts. The first captures innovation by incumbents of type ϕ_h. As is shown below, a firm's innovation effort is a function of its intangible cost parameter and product count, which explains the inclusion of the summation over n. The Poisson rate is multiplied by measure $M_n(\phi_h)$ to obtain the innovation rate. The second term measures innovation by entrants. It is the product of the entry rate e and the probability $G(\phi_h)$ that the entrant is of type ϕ_h.

G. Equilibrium

I now characterize the full stationary equilibrium where productivity, output and wages grow at rate g.

Optimal Innovation Decisions

Firms choose the level of spending on research and development that maximizes firm value. The associated value function, where notation is borrowed from Akcigit and Kerr (2018), reads as

$$rV_t(\phi_i,J_i)-\dot{V}_t(\phi_i,J_i)=\max_{x_i}\left\{\begin{array}{c}\sum_{j\in J_i}\left[\pi_t(\phi_i,\lambda_{ij})+\tau(\phi_i)\left[V_t(\phi_i,J_i\setminus\{\lambda_{ij}\})-V_t(\phi_i,J_i)\right]\right]\\ +x_iP(\phi_i)\mathbb{E}_{\phi_i}\left[V_t(\phi_i,J_i\cup\{\lambda_{ih}\})-V_t(\phi_i,J_i)\right]\\ -w_t\eta^x(x_i)^{\psi^x}n_i^{-\sigma}-F(\phi_i,n_i)\end{array}\right\},\tag{16}$$

where r is the interest rate and where $\dot{V}_t$ denotes the change in V_t with time. The top right-hand line contains the sum of all good-specific items. It is the sum of profits (9) and the change in firm value if the firm ceases production of good j due to creative destruction. $V_t(\phi_i,J_i\setminus\{\lambda_{ij}\})$ denotes the value of producing the set of goods J_i except good j with innovation realization λ_{ij}. The bottom two lines are not specific to goods. The first of these gives the expected increase in firm value from innovation. $\mathbb{E}_{\phi_i}V(\phi_i,J_i\cup\{\lambda_{ih}\})$ denotes the firm's value if it successfully takes product h, taking conditional expectations over λ_{ih} for firm type ϕ_i. The change in value is multiplied by the innovation rate and the probability $P(\phi_i)$ that the firm is able to offer a sufficiently low quality-adjusted price.[20] The final line gives the costs of R&D and a fixed term $F(\phi_i,n_i)$. Firms must pay the latter in order to operate, and it is assumed to equal the option value of R&D. This ad-hoc restriction, borrowed from Akcigit and Kerr (2018), ensures that the value function is linear in the number of goods that firms produce, so that the model admits an analytical first-order condition. In Section V I remove this assumption and show that, though significantly reducing tractability, the results are qualitatively and quantitatively robust.

Proposition 1. *The value function of a firm with intangible cost ϕ_i that produces a portfolio of goods J_i with cardinality n_i grows at rate g along the balanced growth path and is given by*

$$V(\phi_i,J_i)=\sum_{j\in J_i}\pi(\phi_i,\lambda_{ij})(r-g+\tau(\phi_i))^{-1},$$

which is decreasing in ϕ_i. The optimal rate of innovation reads as

$$x_{n_i}(\phi_i)=\left(P(\phi_i)\mathbb{E}_{\phi_i}\left[\frac{\pi_t(\phi_i,\lambda_{ij})}{r-g+\tau(\phi_i)}\right](\eta^x\psi^xw)^{-1}\right)^{\frac{1}{\psi^x-1}}n_i^{\frac{\sigma}{\psi^x-1}}.\tag{17}$$

The optimal entry rate is given by

$$e=\left(\sum_{\phi_e\in\Phi}G(\phi_e)P(\phi_e)\mathbb{E}_{\phi_h}\left[\frac{\pi(\phi_e,\lambda_{ej})}{r-g+\tau(\phi_e)}\right](\eta^e\psi^ew)^{-1}\right)^{\frac{1}{\psi^e-1}}.\tag{18}$$

Proof: Appendix A.

[20] Formally, this probability is given by

$$P(\phi_i)=\sum_{\phi_{\tilde{i}}\in\Phi}K(\phi_{\tilde{i}})H\left(\frac{p^c(\phi_i)}{p^c(\phi_{\tilde{i}})}\right),$$

where $K(\phi_{\tilde{i}})$ is the fraction of products in the economy that is produced by firms with an intangible cost parameter of $\phi_{\tilde{i}}$. Once a firm has successfully innovated, its profits are determined solely by innovation step size λ_{ij}, and hence the value of an additional product does not depend on the identify of the previous producer.

First-order condition (17) is intuitive. Firms engage in more innovation when the expected increase in value is higher, and invest less when the innovation cost-parameters are high. Innovation increases in the firm-size n_i — although if $\sigma < \psi^x - 1$, the firm's expected growth rate will decline with size. Firms with a lower intangible cost parameter ϕ_i choose a higher innovation rate because their ability to reduce marginal costs increases profitability. They furthermore face a lower rate of creative destruction, which decreases the effective discount rate. Firms with a lower ϕ_i also have a higher probability of successfully becoming the new producer of products upon which they innovate. Jointly, these effects cause a negative relationship between ϕ_i and the rate of innovation.

Innovation by entrants (18) is such that the marginal cost of increasing the entry rate e is equal to the expected value of producing a single good, adjusted for the probability that the entrant is able to take over production from the incumbent by offering a sufficiently low quality-adjusted price. Because entrants only learn about their type after they have drawn an innovation, the expectation of the value of producing a good is taken over the distribution of firm types at entry $G(\phi)$.

Dynamic Optimization by Households

Maximizing life-time utility with respect to consumption and savings gives the usual Euler equation,

$$\frac{\dot{C}}{C} = r - \rho, \tag{19}$$

combined with the transversality condition. Along the balanced growth path, consumption grows at the same rate as output and productivity, so that $r - g = \rho$.

Firm Measure and Size Distribution

The optimal innovation rate in (17) is a function of a firm's intangible input costs ϕ_i and the number of goods n_i it produces. The rate of creative destruction (and hence the growth rate of output and productivity) therefore depends on the equilibrium distribution of n and ϕ across firms. Along the balanced growth path, these distributions are stationary. To find the stationary distributions, consider the law of motion for the measure of firms that produce more than one product:

$$\dot{M}_n(\phi_i) = \left(M_{n-1}(\phi_i)x_{n-1}(\phi_i) - M_n(\phi_i)x_n(\phi_i)\right) P(\phi_i) + \left(M_{n+1}(\phi_i)[n+1] - M_n(\phi_i)n\right)\tau(\phi_i), \tag{20}$$

where the first term captures entry into and exit out of measure $M_n(\phi_i)$ through innovation by firms of type ϕ_i with $n-1$ products and n products, respectively. The second term captures entry and exit of firms with $n+1$ and n products that ceased producing one of their products through creative destruction. For the measure of single-product firms, the law of motion reads as

$$\dot{M}_1(\phi_i) = \left(eG(\phi_i) - x_1(\phi_i)M_1(\phi_i)\right) P(\phi_i) + \left(2M_2(\phi_i) - M_1(\phi_i)\right)\tau(\phi_i). \tag{21}$$

The stationary firm-size distribution follows from setting both equations to zero for each n.

The fraction of goods that is produced by firms with intangible cost ϕ_i is given by

$$K(\lambda_i) = \frac{\sum_{n=1}^{\infty} n M_n(\lambda_i)}{\sum_{\lambda_h \in \Phi} \sum_{n=1}^{\infty} n M_n(\lambda_h)}. \tag{22}$$

Labor Market Equilibrium

The solutions to the static and dynamic optimization problems of firms allow the labor market equilibrium conditions to be defined. Labor is supplied inelastically by households at a measure standardized to 1. Equilibrium on the labor market requires that

$$1 = L^p + L^f + L^{rd} + L^e,$$

where L^p is the labor used to produce intermediate goods. Inserting the unit-elastic demand function, markup (8) and intangible first-order condition (7) into $L^p = \int_0^1 \mathbf{1}_{j \in J_i} l_{ij} \, di \, dj$ yields

$$L^p = \int_0^1 \int \mathbf{1}_{j \in J_i} \frac{Y}{w} \left[1 - \left(\kappa_{ij} \frac{w}{Y} \theta \lambda_i \right)^{\frac{1}{\theta+1}} \right] \kappa_{ij}^{-1} \, di \, dj,$$

where $\mathbf{1}_{j \in J_i}$ is the indicator function that equals one when firm i produces good j. L^f is the labor used to fulfill the intangible fixed costs:

$$L^f = \int_0^1 \int \mathbf{1}_{j \in J_i} \left[\left(\kappa_{ij} \frac{w}{Y} \theta \lambda_i \right)^{-\frac{\theta}{\theta+1}} - 1 \right] \lambda_i \, di \, dj.$$

L^{rd} is the labor involved with research and development carried out by existing firms:

$$L^{rd} = \sum_{\lambda_i \in \Phi} \sum_{n=1}^{\infty} \left[M_n(\lambda_i) \eta^x x_n(\lambda_i)^{\psi^x} n^{-\sigma} \right],$$

while L^e is the labor involved with research and development carried out by entrants $L^e = \eta^e e^{\psi^e}$, where innovation rates $x_n(\lambda_i)$ and e are dynamically optimized in line with (17) and (18).

Aggregate Variables

I can now characterize the economy's aggregate variables. The equilibrium wage is given by

$$w = \exp\left(\int_0^1 \int \mathbf{1}_{j \in J_i} \ln \left[\frac{q_{ij}}{s_{ij}} \right] di \, dj \right) \exp\left(\int_0^1 \int \mathbf{1}_{j \in J_i} \ln \left[\frac{s_{ij}}{\kappa_{ij}} \right] di \, dj \right). \tag{23}$$

The first term of (23) is the standard CES productivity term. The second term is the inverse of the expected markup. Note that a rise in the intangibles has no effect on the level of the wage because s_{ij} cancels out. While a firm that deploys more intangibles becomes productive, it is able to proportionally raise its markups. These have offsetting effects on the level of the wage. Aggregate output is given by

$$Y = L^p \exp\left(\int_0^1 \int \mathbf{1}_{j \in J_i} \ln \left[\frac{q_{ij}}{s_{ij}} \right] di \, dj \right) \frac{\exp \int_0^1 \int \mathbf{1}_{j \in J_i} \ln \kappa_{ij}^{-1} \, di \, dj}{\int_0^1 \int \mathbf{1}_{j \in J_i} \kappa_{ij}^{-1} \, di \, dj}. \tag{24}$$

Derivations are provided in Appendix A.3. As in the model with heterogeneous markups and misallocation by Peters (2020), the last term captures the loss of efficiency due to the dispersion of markups. If all markups are equalized, the term is equal to 1, while it declines as the variance of markups increases. Total factor productivity is the product of the second- and the last term in (24).

Equation (24) reveals that a rise in intangibles has two counteractive effects on the level of output. The spread of markups increases when the average s_{ij} falls given (8), because a lower s_{ij} amplifies the heterogeneity in markups caused by the heterogeneous innovation steps (the second term in (24)). On the other hand, the decrease in s_{ij} has a direct positive effect on productivity because it increases the CES productivity index (the first term in (24)). As will be clear below, the second effect dominates in feasible calibrations. That means that a rise in intangibles initially has a positive effect on the level of output and on productivity. The next section shows, however, that this may not be the case for growth.

Growth

The growth rate of total factor productivity and output is a function of creative destruction.

Proposition 2. *The constant growth rate of productivity, consumption C, output Y and wages w is*

$$g = \sum_{\psi_i \in \Psi} K(\psi_i)\tau(\psi_i)\mathbb{E}_{-\psi_i}(\lambda_{hj} - 1), \tag{25}$$

where $\mathbb{E}_{-\psi_i}(\lambda_{hj})$ is the expected realization of λ_{hj} when a firm with ψ_i is the incumbent of a good before a different firm h becomes the new producer due to successful innovation. **Proof:** Appendix A.

The proposition shows that growth is the expected increase in quality multiplied by the rate of creative destruction, weighted by the fraction of product lines that firms of each intangible cost own.

Equation (25) shows the counteracting effects of an increase in ψ_i at a subset of firms. On the one hand, firms with a lower ψ_i have a greater incentive to invest in research and development, which raises the rate of creative destruction. On the other hand, even at a constant innovation rate, the presence of low-ψ_i firms has a negative effect on the rate of creative destruction because high-ψ_i firms have a lower probability of becoming the new producer. This has not only a direct effect on growth at given innovation rates, but also an indirect effect, as these firms reduce their expenditure on R&D.

Equilibrium Definition

Definition 1. *The economy is in a balanced growth path equilibrium if for every t and for every intangible productivity $\psi_i \in \Psi$, the variables $\{r, e, L^p, g\}$ and functions $\{x_{n_i}(\psi_i), K(\psi_i), M_{n_i}(\psi_i), s^*_{ij}, \tau(\psi_i)\}$ are constant, $\{Y, C, w,\}$ grow at a constant rate g that satisfies (25), interest rates follow from (19), Y satisfies (24), innovation rates $x_{n_i}(\psi_i)$ satisfy (17), the entry rate e satisfies (18), firm distribution K_{ψ_i} and measure M_{ψ_i} are constant and satisfy (20) and (21), markups μ_{ij} satisfy (8), the fraction of marginal costs reduced through intangibles s^*_{ij} satisfies (7) for all λ_{ij}, the rate of creative destruction $\tau(\psi_i)$ satisfies (14), and both goods and labor markets are in equilibrium so that $Y = C$ and $L^p = 1 - L^s - L^{rd} - L^e$.*

II. Model Meets Data

This section empirically validates the main mechanisms of the model. Section A introduces firm-level micro-data for two countries, the United States and France. Section B derives predictions that are tested with this data, using a new fixed costs measure detailed in Section C. Evidence on the empirical relationship between fixed costs, intangibles, innovation and markups is presented in Section D.

A. Data

To perform the empirical analysis for the United States I use micro data from financial statements on listed firms, while for France I use administrative tax data on the universe of firms. Online Appendix D, replicating the macroeconomic trends that motivate this paper for France, confirms that the country has incurred a decline in productivity growth and business dynamism, and a modest rise in markups.

Data for U.S. firms comes from S&P's Compustat (Standard and Poor's 2019). Compustat contains balance sheet and income statement data for all listed firms in the U.S. I restrict the sample to firms outside of finance, insurance and real estate between 1979 and 2015, and drop firms with missing or negative sales, (fixed) assets and operating expenses. The sample covers 11,750 firms across 788 6-digit NAICS industries.

The French data comes from two administrative datasets on the universe of firms outside of finance and agriculture (FICUS, from 1994 to 2007, and FARE, from 2008 to 2016), both based on data from the tax office DGFiP (Insee 2019). FICUS and FARE share an identifier (the *siren code*) that consistently tracks firms over time. The data contains the full balance sheet and income statement, with detailed breakdowns of revenues and costs. The sample covers 1,087,726 firms across 651 NACE industries.[21]

Details on variable definitions and data construction are provided in Online Appendix B. Summary statistics are provided in Table 1. All variables are deflated and are winsorized at their 1% tails.

B. Testable Predictions

These micro datasets enable a test of the model's main predictions. Before deriving the predictions, I generalize production function (2) to facilitate a mapping to the data. Assumptions on the remainder of the framework follow Section I. I generalize (2) by assuming that firms produce with a first-degree homogeneous production function $z(z_{ijt,1}, z_{ijt,2}, .., z_{ijt,k})$ with k tangible production factors or intermediate inputs. Intangibles allow firms to reduce marginal costs, in exchange for higher fixed costs. Denoting by s_{ijt} the share of marginal costs that a firm keeps, the production function reads

$$y_{ijt} = s_{ijt}^{-1} z(z_{ijt,1}, z_{ijt,2}, .., z_{ijt,k}). \tag{26}$$

The first-degree homogeneity of $z(\cdot)$ implies that firm i's marginal cost is $mc_{ijt} = s_{ijt} c(w_{1t}, w_{2t}, .., w_{kt})$, where w_{kt} denotes the factor price of tangible production factor k at time t.

[21] The FICUS-FARE panel was created for Burstein, Carvalho and Grassi (2019), who kindly allowed me to use their data in this paper.

Table 1: Descriptive Statistics

	Mean	Std. Dev.	Median	10th Pct.	90th Pct.	Obs.
U.S. Compustat Firms (1979-2015)						
Sales (revenue)	1,733,510	5,154,239	1,94,318.5	13,329.7	3,603,252	127,682
Operating expenses	1,466,022	4,326,827	169,248.4	12,957.5	3,074,281	127,682
Cost of goods sold	1,111,563	3,315,834	116,819.2	7,133.9	2,351,005	127,682
Selling, General, and Adm. expenses	335,829.9	1,046,907	39,864.5	3,918.3	650,1041	127,682
Capital stock	1,109,350	3,714,094	79,352.3	4,910.0	2,131,897	127,682
French Firms in FICUS-FARE (1994-2016)						
Sales (revenue)	4,684	103,282	617	149	4,996	9,913,058
Employment (headcount)	19	356	5	1	28	9,913,058
Wage bill	622	10,753	144	37	831	9,913,058
Capital stock	1,738	131,183	92	12	895	9,913,058
Intermediate inputs and raw materials	2,235	58,700	136	0	1,924	9,913,058
Other operating expenses	1,211	35,656	124	33	1168	9,913,058

Notes: Nominal figures in thousands of deflated dollars (U.S.) and euros (France). Sales, operating costs and materials are deflated with KLEMS sector deflators; the wage bill and capital are deflated with the GDP deflator.

The generalized model yields a number of empirical predictions. First, one of the core assumptions of the model is that intangible inputs are fixed costs. It follows directly that if one observes both (a subset of) a firm's intangibles and its fixed costs, these should correlate positively. Second, the model yields equilibrium relationships between intangibles, research and development expenditures and markups, as summarized by the following proposition:

Proposition 3. *For equilibria where firms choose a positive level of spending on intangibles, the model implies the following firm-level relationships between intangibles, markups and innovation:*

a. The cost-minimizing fixed- over total cost ratio decreases in a firm's intangible cost parameter $\acute{A}_i$.

b. Firms with lower intangible cost parameters $\acute{A}_i$ have higher spending on research and development, and have higher markups across their products than other firms do.

Proof: Appendix A.

Combined, parts (a) and (b) of the proposition invite regressions that relate firms' ratios of fixed costs over total costs to research and development and markups, with a positive predicted correlation.

C. Measuring Fixed Costs

Testing these predictions involves two main challenges. First, firms do not identify whether costs are variable or fixed in financial statements. Second, the administrative datasets provide data at the firm level, while production function (26) operates at the product level. I address both issues below.

Measure

I derive a new time-varying measure of fixed costs from the difference between marginal cost markups and the profit rate, which equals operating profits over revenue. Past work typically uses the sensitivity of a firm's operating costs or profits to sales shocks to measure fixed costs, under the assumption that all variable costs are set freely.[22] This is problematic when firms face adjustment costs for some variable inputs (e.g. when adjusting their labor force), and it does not yield a time-varying measure of fixed costs at the firm level. As posts on the income statement are broad, they too cannot be classified as fixed costs, even in the French data. The new measure instead takes advantage of marginal costs being constant within firm-product-years, implying that the accounting definition of the profit rate is

$$\frac{\pi_{ijt}}{p_{ijt}y_{ijt}} = \frac{\left(p_{ijt} - mc_{ijt}\right) y_{ijt}}{p_{ijt}y_{ijt}} - \frac{f(s_{ijt},\psi_i)}{p_{ijt}y_{ijt}},$$

which follows from dividing profits (9) by revenue. Isolating fixed costs on the left, and defining the markup μ_{ijt} for the ratio of prices to marginal costs, yields a model-consistent measure of fixed costs:

$$\frac{f(s_{ijt},\psi_i)}{p_{ijt}y_{ijt}} = \left(1 - \frac{1}{\mu_{ijt}}\right) - \frac{\pi_{ijt}}{p_{ijt}y_{ijt}}. \tag{27}$$

The equation shows that profits differ from markups because the latter measure *marginal* profitability, while the profit rate measures *average* profitability. A firm with positive fixed costs should have a profit rate below the markup. This implies that rising markups do not necessarily reflect rising profitability.

It is straightforward to aggregate the product-level (27) to a model-consistent measure of fixed costs at the firm level. A firm-level measure is needed because the micro datasets contain firm-level revenues $py_{it} = \sum_{j\in J_{it}} p_{ijt}y_{ijt}$ and operating profits $\pi_{it} = py_{it} - tc_{it}$, where tc_{it} denotes firms' total costs and where J_{it} is a firm's product portfolio. Because firms may face fixed costs that are not related to intangible spending on individual products (such as overhead), I write total costs tc_{it} as follows:

$$tc_{it} = \tilde{f}_{it} + \sum_{j\in J_{it}} \left(y_{ijt}s_{ijt}\mathbf{c}(w_{1t}, w_{2t}, .., w_{kt}) + f(s_{ijt},\psi_i)\right),$$

where $\tilde{f}_{it}$ are the additional firm-level fixed costs that are unrelated to intangibles and that apply at the firm level. The following proposition applies:

Proposition 4. *A firm's total fixed costs* $f_{it} = \tilde{f}_{it} + \sum_{j\in J_{it}} f(s_{ijt},\psi_i)$ *are identified by*

$$\frac{f_{it}}{py_{it}} = \left(1 - \frac{1}{\mu_{it}}\right) - \frac{\pi_{it}}{py_{it}}, \tag{28}$$

where μ_{it} *is the harmonic average of the firm's product-level markups* μ_{ijt}. **Proof:** Appendix A.

It follows that firm-level fixed costs can be identified from firm-level data on operating profits and revenue (both observable on financial statements), as well as markups. Markups are not observable because financial statements lack data on marginal costs and prices. I therefore estimate markups using the method proposed by Hall (1988). He shows that under cost minimization, a firm's markup is given by the product of the output elasticity of a variable input multiplied by the ratio of a firm's sales to its expen-

[22] Online Appendix C shows that trends in alternative measures of fixed costs are similar to the trend in the new measure.

diture on that input. Online Appendix A.8 derives that, if a firm's products have equal output elasticities, Hall's firm-level markup estimates are the harmonic average of product-level markups.

The advantage of the Hall (1988) methodology to estimate markups is that it does not assume a particular market structure or competition, and that it is consistent with the framework in Section I. The measure requires firm-level data on revenue and spending on some variable input, which is assumed to be set flexibly and without adjustment costs. Other inputs may be fixed, variable or a combination of both: as long as one freely-set variable input is observed, markups can be estimated consistently. For the French data I use intermediate inputs as the variable input, while for U.S. data I use markups from De Loecker, Eeckhout and Unger (2020), which use cost of goods sold. Revenue and expenditure on the input are observed on the income statement, while I obtain the output elasticity by estimating a translog production function using the procedure proposed by De Loecker and Warzynski (2012).[23]

A final data issue is the mapping of Proposition 3 from the product to the firm level in the presence of firm overhead costs $\tilde{f}_{it}$. When firms produce additional goods, the fixed-cost ratio declines mechanically because firm-level $\tilde{f}_{it}$ is now spread across more products. As firm concentration has increased over time in the data, an analysis of aggregate trends in fixed costs underestimates the rise in fixed costs that is driven by fixed-cost intangibles. In the regressions I address this by controlling for (log) revenue, as revenue is proportional to firm size ($py_{it} \;/\; n_{it}$) in the model. Revenue, therefore, controls for the mechanical decline in the ratio of $\tilde{f}_{it}$ over total costs in firm size.

Fixed Cost Estimates

I use the data on operating expenses, revenue and markups to measure firm-level fixed costs for France and the United States following (28). On average, fixed costs comprise 23.7% of operating costs in the U.S. data and 19.4% in the French data, with respective standard deviations of 19.6% and 22.6%.

Figure 2 plots the revenue-weighted average ratio of fixed to total costs. I weigh by revenue, because revenue is proportional to the firm's weight in aggregate fixed costs in the model. The measure shows a persistent increase in both countries. Fixed costs made up less than 13% (9.5%) of costs for U.S. (French) firms at the start of the sample, and close to 23% (14%) at the end. Over the full episode, the increase is greater for U.S. firms but this seems to be due to the difference in time samples. Between 1995 and 2015, the increase in the fixed-costs share is around 5 percentage points in both datasets.

Online Appendix C provides various robustness checks for the trend in fixed costs. It shows that the results are robust to alternative estimates of markups, as well as to various approaches to the adjustment of profits for capital costs.[24] The appendix also confirms that alternative measures of fixed costs from the literature yield similar levels and trends as those in Figure 2. Finally, the appendix also contains an illustration of the sectoral composition of fixed costs (Figure A4). It shows that fixed costs are especially high in the information sector, while variable costs are relatively important in retail and wholesale.

[23] Details are provided in Online Appendix C.

[24] The definition of operating profits for Figure 2 accounts for depreciation costs, but it does not adjust profits for the rental costs of capital. The appendix derives multiple series of required rental costs r_t^K from the estimates of risk-free interest rates and capital risk premia in Caballero, Farhi and Gourinchas (2017), and shows that the trends in Figure 2 are robust to the adjustment.

Figure 2. Weighted-Average Ratio of Fixed Costs to Total Costs

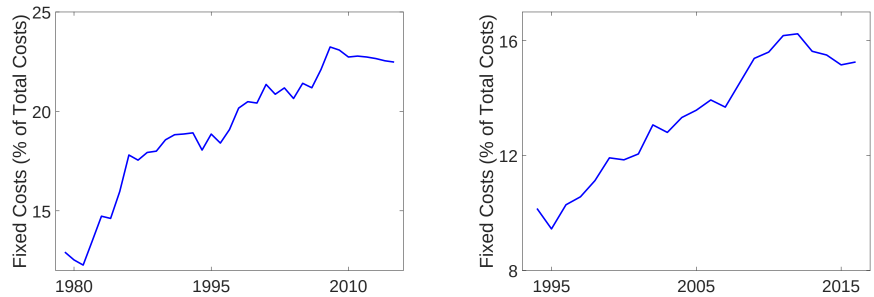

(a) United States (b) France

Notes: Sales-weighted average of fixed costs as a percentage of total costs, U.S. listed firms (left) and universe of French firms (right). Fixed costs are derived from the difference between markups and profit rates at the firm level in line with (28).

Nearly all broad sectors have seen an increase in their share of fixed costs in total costs, and a formal between-within decomposition in Appendix Table A3 confirms that the increase occurs largely within sectors.

A key assumption in the model is that intangible inputs are fixed costs. While there is no exhaustive information on firm-level spending on intangible inputs, external data on software investments is available. As software spending is a subset of total intangible spending, we should see it positively correlate with fixed costs. To assess whether this is the case in the data, I estimate

$$\ln\left(\frac{f_{it}}{py_{it}}\right) = \beta \ln\left(\frac{f^s_{it}}{py_{it}}\right) + \gamma' x_{it} + \epsilon_{it}, \tag{29}$$

where software spending is denoted by f^s_{it}, so that $\beta > 0$ if f^s_{it} is indeed part of the firm's fixed costs $\sum_{j \in J_{it}} f(s_{ijt}, \phi_i)$. Both the dependent and the explanatory variable are standardized by revenue because revenue is proportional to a firm's product count in the model. This controls for the mechanical correlation between fixed costs and software spending that arises when firms produce more goods.[25]

To measure software for U.S. firms, I use estimates of annual software budgets in the Ci Technology Database (CiTDB), produced by the marketing company (Harte-Hanks 2017). The CiTDB collects site-level IT data through phone surveys, which is sold for the purpose of commercial acquisition. It contains consistently defined estimates of software budgets for 2010 and for 2012-2015.[26] Appendix B details the data construction. The dataset contains 6,585 observations.

For France, data on software and IT comes from the *Enquête Annuelle d'Entreprises* (EAE, Insee 2019), which is an annual survey that (post-weighting) representatively samples around 12,000 firms between 1994 and 2007. The survey provides a comprehensive panel of firms with more than 20 employees,

[25] Alternatively, the regression could be performed with *controls* for the log of revenue. This yields similar results.

[26] I thank Nick Bloom and John Van Reenen for graciously providing an extract from the CiTDB to perform this analysis. I additionally observe firm-level adoption of personal computers. While personal computers are not directly a measure of intangible inputs, they are a common proxy for IT intensity at the firm level (see, e.g., Bloom, Sadun and Van Reenen 2012), and this data is available for a longer sample from 1997 to 2015. The dataset also contains detailed data on the kind of IT systems that firms have installed, but it is not possible to distinguish between missing entries and entries where a system was not installed.

Table 2: Relationship between Software Spending and Fixed-Cost Share

Fixed-Cost Share (log)	I	II	III	IV	V	VI
United States (CiTDB, 2010 to 2015)						
Software budget over revenue (log)	0.093	0.067	0.054	0.055	0.022	0.019
	(0.008)	(0.010)	(0.010)	(0.010)	(0.007)	(0.007)
R^2	0.054	0.063	0.170	0.172	0.015	0.038
France (EAE, 1994 to 2007)						
Software investments over revenue (log)	0.130	0.117	0.0719	0.0678	0.0456	0.0324
	(0.002)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
R^2	0.039	0.056	0.192	0.197	0.032	0.129
Year fixed effects	No	No	No	Yes	No	Yes
Firm fixed effects	No	No	No	No	Yes	Yes
Industry fixed effects	No	No	Yes	Yes	No	No
Revenue (product-count) control	No	Yes	Yes	Yes	Yes	Yes

Notes: Dependent variable is fixed costs over revenue (log). Explanatory variable is (log) software investments over revenue. Revenue is deflated with the sector-specific gross output deflator, software with the input deflator from EU-KLEMS. Variables are winsorized at 1% tails. Firm-clustered standard errors are given in parentheses. Revenue in logs, as it is proportional to a firm's product count in the model. Sector fixed effects are at the 2-digit level. U.S. regressions have 6,635 observations; French regressions have 106,865 observations.

and samples smaller firms in most sectors. I use this survey for data on software investments, developed in-house or purchased externally. The matched data contains 106,865 firm-years across 31,005 firms.

Results are presented in Table 2. Column I presents univariate regressions, which show a significantly positive relationship between software spending and fixed costs in both countries. Column II adds log revenue to control for firm size n_{it}. Adding 2-digit industry- and year fixed effects in columns III and IV has limited effects on the estimated correlations, which are comparable in size for France and the U.S. The coefficients shrink substantially when adding firm fixed effects. This is consistent with the model, as most of the variation in intangibles comes from Á_i, which is fixed within the firm.

Results in Table 2 are robust to alternative specifications, such as controlling for size through costs instead of revenue as a control, or to a broader measure of information technology budgets to measure software spending in the U.S. data. Appendix Table A5 shows that there is also a positive correlation between fixed costs and personal computer intensity. I also find a significantly positive relationship when f^s_{it} is proxied with dummies for the adoption of specific technologies.[27] I conclude that there is a robustly positive relationship between observed intangible inputs and fixed costs, in line with the model.

D. Correlations between Fixed Costs, Research and Development, and Markups

To further explore the empirical validity of the model's mechanisms, I now assess the firm-level conditional correlations implied by the model's equilibrium (Proposition 3) in the micro data.

[27] Data on technology adoption is available for France. I use the Enquête sur les Technologies de l'Information de la Communication (TIC, Insee (2019)), which details a firm's installed IT systems from 2008 to 2016. It samples around 10,000 firms annually. Online Appendix F's Table A4 regresses (log) fixed costs over revenue on IT system adoption and finds a strong relationship.

Table 3: Relationship between Fixed-Cost Share and Research & Development

R&D intensity (log)	I	II	III	IV	V	VI
United States						
Fixed costs over total costs (log)	0.693	0.623	0.566	0.536	0.179	0.163
	(0.018)	(0.018)	(0.018)	(0.018)	(0.012)	(0.012)
R^2	0.175	0.231	0.306	0.320	0.028	0.043
France						
Fixed costs over total costs (log)	0.412	0.336	0.123	0.128	0.044	0.053
	(.017)	(0.017)	(0.015)	(0.015)	(0.024)	(0.025)
R^2	0.045	0.103	0.275	0.278	0.015	0.026
Year fixed effects	No	No	No	Yes	No	Yes
Firm fixed effects	No	No	No	No	Yes	Yes
Industry fixed effects	No	No	Yes	Yes	No	No
Revenue (product-count) control	No	Yes	Yes	Yes	Yes	Yes

Notes: Dependent variable is R&D intensity (log). Explanatory variable is fixed costs over total costs (log). Variables are winsorized at 1% tails. Firm-clustered standard errors are given in parentheses. Revenue in logs, as it is proportional to a firm's product count in the model. Sector fixed effects are at the 2-digit level. U.S. regressions have 58,246 observations; French regressions have 20,666 observations.

Research and Development

I first assess the relationship between fixed costs and innovative investments, measured through R&D intensity. R&D intensity is measured as the ratio of R&D over revenue, as is standard in the literature (e.g. Hall, Mairesse and Mohnen 2010). In equation (A.8) of Appendix A I show that the firm's R&D intensity is a function of three terms. The first term captures the value of becoming the producer of an additional good, which is higher for firms with low intangible costs ϕ_i. Along the balanced growth path, the term is entirely captured by a firm fixed effect. The second term captures that innovation intensity falls in firm size. The final term is a time fixed effect. The firm's intangible costs ϕ_i are unobservable, but Proposition 3 shows that high-intangible firms have higher ratios of fixed costs over total costs. I therefore estimate

$$\ln\left(xrd_{it}/py_{it}\right) = \beta \ln\left(f_{it}/tc_{it}\right) + \gamma \ln(py_{it}) + \delta_t + \varepsilon_{it}, \tag{30}$$

where δ_t are time fixed effects and ε_{it} are residuals, while revenue is included to control for firm size.

I measure xrd_{it} as R&D on the income statement in Compustat, which is feasible because R&D is expensed under U.S. accounting rules.[28] For France I obtain R&D from the *Enquête Communautaire sur L'Innovation* (CIS). The CIS is an innovation survey that was conducted in 1996 and 2000, and biannually since 2004. The variable for xrd_{it} that I take from this dataset is expenditures on innovation activities, including externally purchased R&D and expenditures on external knowledge or innovation-related capital expenditures. The dataset contains 20,666 firm-years across 12,879 firms.

Table 3 presents the results. Estimates in column (I) come from a univariate regression, which shows that firms with relatively high fixed costs indeed have higher R&D intensities than other firms. Controlling for firm size through revenue slightly lowers the coefficient, in line with the prediction of the model that both R&D intensity and fixed costs over total costs (through overhead $\tilde{f}_{it}$) decline in firm size. The

[28] Because R&D is expensed, it negatively affects profits. When calculating fixed costs, I do not include R&D in the definition of costs, thereby preventing a mechanically positive relationship between fixed costs and R&D.

estimates imply that a one percent increase in fixed over total costs raises R&D intensity by 0.3 to 0.6 percent, on average. This result is robust to industry- and year fixed effects.

The model predicts that, in the steady state, fixed costs are orthogonal to innovation rates after conditioning on firm effects. Columns V and VI show that this is not the case in the data, although the coefficients are significantly smaller than in the columns with fewer controls. As Table 3 uses the full sample, it is unlikely that the economy is in the steady state, explaining the positive estimates.[29]

Markups

I next relate fixed costs to markups. As shown in Proposition 3, fixed costs should correlate positively with markups. In a regression, two counteracting forces govern this relationship. To see this, note that product-level markups in the model are $\ln \mu_{ijt} = \ln \lambda_{ij} - \ln s_{ijt}$, where λ_{ij} denotes the innovation step size and where time subscripts were added to s_{ijt}. Given Cobb-Douglas demand, this equals

$$\ln \mu_{ijt} = -\ln s_{ijt} + \ln Y_t - \ln y_{ijt} - \ln w_t.$$

The first counteracting force is captured by the final term, as a higher s_{ijt} directly translates to lower fixed costs. This creates the positive simple correlation between markups and fixed costs. The second force operates through y_{ijt} and λ_{ij}: a greater step λ_{ij} raises markups and lowers fixed costs, because higher markups reduce demand. This creates a negative correlation between markups and fixed costs. Indeed, solving for λ_{ij} in terms of s_{ijt} using first-order condition (7), we can denote markups as

$$\ln \mu_{ijt} = \theta \ln s_{ijt} - \ln \phi_i + \ln Y_t - \ln w_t - \ln \theta.$$

It follows that the relation between fixed costs and markups should be positive in a regression without fixed effects. Conversely, given that fixed costs are higher when s_{ijt} is low, the relationship could be negative when controlling for firm fixed effects.[30]

To test these predictions, I run a firm-level regression of the (log) markup on the (log) ratio of fixed costs over total costs. Compared to the other regressions there is an added complication: fixed costs are derived from markups, so that measurement error in markups causes a positive correlation. I address this by instrumenting total fixed costs with the observable subset of fixed costs: software spending.[31]

Table 4 presents the results. As expected, OLS results in columns I and II show a strong positive correlation between markups and fixed costs, regardless of controls. When addressing bias through correlated measurement error in 2SLS results (columns III and IV), we see that the correlation is strongest in the univariate regressions. Those estimates are economically significant: a one-percent increase in fixed costs raises markups by 0.3 to 0.4 percent. Consistent with the model, fixed effects reduce the

[29] Results in Table 3 are robust to various alternative specifications, including the use of fixed costs over revenue as the explanatory variable. For the U.S. it is also feasible to use software investments from the CiTDB (Online Appendix Table A6). For France this is not feasible, because the R&D and software data comes from different surveys with insufficient overlap.

[30] Firm fixed effects orthogonalize fixed costs from ϕ_i up to the first order. In the model, the relationship between ϕ_i and s_{ijt} is non-linear, so that a linear regression of markups on fixed costs with firm effects may still yield a positive coefficient.

[31] The instrument also controls for measurement error in f_{it} caused by the firm-level overhead $\tilde{f}_{it}$. Controlling for $\tilde{f}_{it}$ by adding the log of revenue as an additional control, as in Table 3, does not affect the qualitative results in Table 4.

Table 4: Relationship between Fixed-Cost Share and Markups

	OLS		2SLS		Reduced Form	
Markup (log)	I	II	III	IV	V	VI
United States (CiTDB, 2010 to 2015)						
Fixed costs over total costs (log)	0.271	0.061	0.439	-2.42	0.043	0.002
	(0.009)	(0.007)	(0.026)	(9.58)	(0.003)	(0.002)
R^2 (First-stage F.)	0.504	0.073	(392)	(9.51)	0.074	0.002
France (EAE, 1994 to 2007)						
Fixed costs over total costs (log)	0.173	0.071	0.307	0.254	0.030	0.003
	(.001)	(0.001)	(0.007)	(0.036)	(0.001)	(0.000)
R^2 (First-stage F.)	0.457	0.449	(1350)	(106.8)	0.032	0.051
Year fixed effects	No	Yes	No	Yes	No	Yes
Firm fixed effects	No	Yes	No	Yes	No	Yes

Notes: Dependent variable is the markup (log). Explanatory variable is fixed costs over revenue (log). Variables are winsorized at 1% tails. Firm-clustered standard errors are given in parentheses. The instrument for the 2SLS regressions is the ratio of software spending over sales. U.S. regressions have 6,635 observations; French regressions have 106,795 observations.

relationship (France) or even reverse its direction (U.S.). The reduced-form regressions confirm that software-intensive firms have higher markups, but not when account is taken of fixed effects.

In summary, the model implies positive correlations between fixed costs and intangibles, markups and innovation, especially when no account is taken of firm fixed effects. The correlations in the micro data are broadly in line with these patterns. I next use the data to quantify the model.

III. Quantification

This section outlines how the model is quantified. I first discuss the calibration and structural estimation strategy, and then discuss the extent to which the model is able to replicate a set of targeted and untargeted moments along the original balanced growth path.

A. Calibration

In the baseline calibration firms have equal intangible costs Á, leaving nine parameters to be calibrated: discount rate ½, the cost elasticity of intangibles µ, innovation cost curvatures Ã^x and Ã^e, innovation cost scalars ´^{x} and ´^{e}, the average innovation step size ¯¸, the parameter governing the firm-size and -growth relationship ¾, and Á. Five parameters are estimated, the others are taken from the literature. I separately estimate the model for the U.S. and France, using Section II's micro data.

Externally Calibrated Parameters

The model is calibrated at an annual frequency. I calibrate the curvature of R&D for entrants (Ã^e) and incumbents (Ã^x) to 2. This is a key parameter because it determines the concavity of the return to R&D. If innovative activities concentrate among fewer firms, the fact that Ã^x È 1 implies that the average effect of these investments on growth is lower. The literature that studies the elasticity of R&D with respect to the user costs ($^{2}{}_{x,w}$) of such activities nds elasticities around -1.0 for tax credit changes (see, e.g., Bloom, Grif th and Van Reenen 2002 for a review). The parameter $\tilde{A}^{x}$ is related to $^{2}{}_{x,w}$ as follows:

$$\tilde{A}^{x} Æ ¡ \frac{^{2}{}_{x,w} ¡ 1}{^{2}{}_{x,w}},$$

and is therefore set to 2. The same value is used in Akcigit and Kerr (2018) and Acemoglu et al. (2018).

I calibrate the curvature parameter µ of xed cost function (3) to match empirical estimates of the pass-through of marginal costs to markups. To see how these relate, note that the rst-order conditions for markups (8) and intangibles (7) imply an equilibrium log markup

$$\ln {}^{1}{}_{ijt} Æ \ln ¸_{ij} ¡ \frac{1}{µÅ1} \ln \left(¸_{ij} \frac{w_t}{Y_t} µ Á_i \right). \tag{31}$$

The elasticity of marginal costs with respect to wages is (µ Å 2)/(µ Å 1), so that the elasticity of markups with respect to marginal costs at a given level of Y is ¡ (µ Å 2)$^{¡1}$. I set µ to 2, which achieves a pass-through of -25%. Empirical estimates of this elasticity vary. Amiti, Itskhoki and Konings (2019) nd a pass-through of -35% in their main results. In robustness checks on the full sample they nd values between -39% and -25%. For rms with fewer than 100 employees they nd coef cients of -3%. Table A13 in Online Appendix F shows that the results are robust to µ Æ 0.86, yielding a -35% pass-through.

As a further robustness check, I calibrate µ from the relationship between variable costs vc_{ijt} Æ $w_t s_{ijt} y_{ijt}$ and xed costs. From the rst-order condition for intangible inputs, it follows that

$$vc^{¤}_{ijt} Æ µ f(s^{¤}_{ijt}, Á_i) Å µ w_t Á_i.$$

I run a rm-level regression of variable costs on xed costs in Table A12 in Online Appendix F. Estimates of µ range from 0.86 to 1.34, similar to the µs derived from markup pass-through. Appendix Table A13 shows that the main results are robust to calibrating µ using the relationship between xed and variable costs. The distribution of the innovation step size $¸_{ij}$ is Pareto, so that

$$H(¸) Æ ¸^{¡\frac{\bar{¸}}{\bar{¸}¡1}},$$

where the shape parameter is a function of the average innovation step size $\bar{¸}$, which I structurally estimate along with the other parameters in the next section. The Pareto distribution ensures that the quality of innovations follows a power law, in line with empirical evidence (e.g. Harhoff, Scherer and Vopel 2003, Kogan et al. 2017). It also delivers an exponential distribution for log markups, as in Peters (2020). The discount rate ½ is set to 0.01, which gives rise to a 2.3% risk-free rate.

Structurally Estimated Parameters

The remaining ve parameters are estimated using indirect inference by matching moments from either the U.S. Compustat data on listed rms or the French administrative data. The U.S. calibration targets moments around 1980, which is the rst year that rm variables from Compustat can be complemented

Table 5: Overview of Parameters

Parameter	Description	Method	Value (U.S.)	Value (France)
ρ	Discount rate	External	.010	.010
θ	Intangibles cost elasticity	External	2.00	2.00
ψ^x	Cost elasticity of innovation (incumbents)	External	2.00	2.00
ψ^e	Cost elasticity of innovation (entrants)	External	2.00	2.00
η^x	Cost scalar of innovation (incumbents)	Indirect inference	3.36	1.73
η^e	Cost scalar of innovation (entrants)	Indirect inference	2.44	2.29
$\bar{\lambda}$	Average innovation step size	Indirect inference	.060	.061
σ	Relationship firm size and firm growth	Indirect inference	.521	.623
ϕ	Intangible costs	Indirect inference	.215	.279

by administrative data on business dynamism. The French calibration targets moments in the first year of the data (1994), or the first available year for surveys.

The structural estimation chooses the set of parameters that satisfies the objective function:

$$\min \sum_{k=1}^{5} \frac{|\text{model}_k - \text{data}_k|}{(|\text{model}_k| + |\text{data}_k|)0.5}\,\omega_k, \tag{32}$$

where model_i and data_i respectively refer to the simulation and data for moment i with weight ω_i. I solve the model as a fixed point along the algorithm described in Online Appendix E. Using the equilibrium values for innovation and entry rates, the firm-size distribution, rates of creative destruction and aggregate quantities such as the efficiency wedge, wages and output, I simulate the economy for 32,768 firms until the distribution of s_{ij} has converged, and simulate ten more years to collect moments.

The following moments are used for the U.S. calibration. I calibrate the intangible efficiency parameter ϕ to match the 1979 ratio of fixed to variable costs of 12.9% in Section II. Online Appendix A.9 derives that this ratio falls in ϕ. The cost scalar of R&D by entrants (η^e) is estimated by targeting the entry rate of 13.8% for 1980 in the Business Dynamics Statistics. The cost scalar of innovation by existing firms (η^x) is estimated by targeting the average ratio of R&D over sales for firms with positive expenditures in 1980, at 2.5%.Following Akcigit and Kerr (2018), I calibrate the parameter that governs the extent to which R&D scales with size (σ) by targeting the following OLS regression of size on growth:

$$\Delta_i\left(py\right) = \alpha_s + \beta \ln\left(p_i y_i\right) + \varepsilon_i, \tag{33}$$

where the left-hand side is the growth rate of sales using the measure of growth in Davis et al. (2006), while α_s is a sector fixed effect. Akcigit and Kerr (2018) estimate (33) on Census data and find β=-0.035, which implies that a firm with 1% greater sales is expected to grow 0.035% less. I estimate average innovation step size $\bar{\lambda}$ by targeting productivity growth along the balanced growth path of 1.3%, which equals average growth of total factor productivity between 1969 and 1980 in the Fernald series. [32]

The calibration for France relies on French counterparts of the U.S. moments. Intangible cost parameter ϕ is calibrated to match the 1994 ratio of fixed to variable costs of 9.5% in Section II. The cost scalar of research and development by entrants (η^e) is estimated by targeting an entry rate of 10%. This

[32] I measure productivity growth with Fernald's utilization-adjusted series. This series is closest the model because it adjusts for temporary fluctuations in the utilization of labor and capital, as such fluctuations do not occur in the model.

Table 6: Comparison of Theory and Data for Targeted Moments

				United States		France	
Parameter	Moment	Weight	-	Model	Target	Model	Target
¯¸	Long-term growth rate of productivity	1		1.3%	1.3%	1.3%	1.3%
Á	Fixed costs as a fraction of total costs	2		12.9%	12.9%	9.5%	9.5%
¾	Relation between rm growth and size	1		-.035	-.035	-.035	-.035
´ e	Entry rate (fraction of rms age 1 or less)	1		13.2%	13.8%	10%	10%
´ x	Ratio of research and development to sales	1		2.5%	2.5%	2.4%	3.2%

is the fraction of rms that enter the FARE-FICUS dataset for the rst time in 1995, the second year for which data is available and therefore the rst year that entry is observed. The cost scalar of innovation by existing rms (´ x) is estimated by targeting the average ratio of R&D over sales in the CIS for 1996, which is 3.1%. I calibrate ¾ to match the ¯ in (33) using the French micro data for 1994-1995. The estimated ¯ is -0.035, coincidentally the same as for the U.S. I target a productivity growth rate of 1.3%, which is the average TFP growth rate between 1969 and 1994 in the Penn World Tables.

Table 5 presents an overview of the calibrated and estimated parameters. The lower R&D intensity of U.S. rms gives rise to a higher innovation-cost scalar ´ x, while the higher ratio of xed- to variable costs of U.S. rms causes their baseline estimated intangible costs Á to be lower than that of the French rms. The estimated innovation-step size is similar for both countries.

B. Model Properties

A comparison of theoretical and empirical targeted moments is provided in Table 6. The rst column lists the parameter that corresponds most closely to the moment, the second column describes the moment, and the third column summarizes the moment's weight in the structural estimation. All moments receive the same weight except the share of xed costs, which is assigned a weight of two. The model is able to match moments on growth, xed costs, and the relationship between rm growth and rm size precisely for both countries. R&D intensity is also precisely matched for the U.S. calibration. Entry is precisely matched for France, although the model underestimates French R&D intensity.

The rm-size distribution is untargeted. The Cobb-Douglas aggregator implies that a rm's revenue is determined by the number of goods that it produces, which is plotted against data in Figure 3. I rely on the Compustat Segments data for the U.S. to count the number of NAICS industries in which rms operate (Figure 3a).[33] This is the orange-circled line. Results show that U.S. listed rms operate in more sectors than the model predicts. Note that the Compustat segments are an imperfect measure of the number of products that rms produce because rms apply heterogeneous reporting standards on what a segment is. Further, a third of rms do not report their segments at all. The green-squared line plots an alternative distribution of the product count, setting the number of products to one for non-reporting rms. This brings the distribution closer to what is predicted. The difference between the fraction of rms with 2 and 3 (and 3 and 4) products is also accurately predicted. Figure 3b plots the same for France. Data comes from the Enquête Annuelle de Production dans L'Industrie (EAP, Insee 2019). This

[33] The rst year with NAICS segment codes is 1990, which is plotted here. Details are provided in Data Appendix B.

Figure 3. Number of Products by Firm: Theory and Data

(a) United States (b) France

Notes: U.S. data are taken from the Compustat Segments le and count the number of primary NAICS codes that rms report to operate in during 1990. Adjusted segments data assign a segment count of 1 for rms that are not included in the segments le. French data are taken from the Enquête Annuelle de Production dans L'Industrie (manufacturing only, 2009).

dataset is available only for rms in manufacturing and contains identi ers for each product that the rm sells.[34] The gure shows that the distribution of the number of product count is closely matched.

Markups are also untargeted. The model predicts average markups of 1.41 for the U.S. calibration and 1.29 for the French calibration. This overstates the actual markups at the beginning of the sample, which average 1.27 and 1.17 in the U.S. and French data, respectively.[35]

Table A7 in Online Appendix F presents a set of additional untargeted moments. The rst panel analyzes the relationship between size and age. Size is measured as sector-de ated sales, while age is measured as years since creation in France, and as years since entry into Compustat for the U.S. Both are transformed to within-year quartiles indexed from 1 to 4. For the U.S., the model accurately predicts that small rms are more likely to exit and less likely to stop producing a product, but cannot explain the relationship between exit and age. This could be because U.S. exits are calculated in Compustat, which may be for other reasons than rm closure. The model correctly predicts that young rms are smaller than older rms, as they have had less time to accumulate patents. The model also matches for France that young and small rms are more likely to exit and less likely to stop producing a product.

Tables A8 to A11 additionally compare regressions of simulated data between the model and the data. Table A8 estimates the elasticity of xed-over-total costs with respect to xed costs; Table A9 estimates the relationship between xed costs and rm size. Both are well-matched. Table A10 and Table A11 present estimations that relate xed costs to markups and R&D as in Section II. These are not well-matched because all rms have equal intangible ef ciencies, but improve in the next section.

IV. Quantitative Analysis

This section contains the main exercise: a quantitative analysis of the effect of a rise of intangibles on productivity growth, business dynamism and markups. Section A rst outlines how high-intangible

[34] The rst year of the survey is 2009, which is plotted here. Details are provided in Data Appendix B.

[35] Appendix G presents an alternative calibration that targets markups, and shows that results are qualitatively robust.

firms are introduced, while Section B analyzes how they change the balanced growth path. The transition path from the old to the new steady state is discussed in Sections C to E.

A. Introducing Heterogeneous Intangible Efficiency

To model the rise of intangibles, I introduce a group of firms with lower intangible costs than the homogeneous intangible cost in the initial calibration of Section III. Two parameters characterize the introduction of these firms: their cost parameter, $\bar{A}$, and the fraction of entrants that receive it, $G(\bar{A})$. To calibrate $G(\bar{A})$, I target the decline in entry. Entry depends on the share of firms with a higher intangible efficiency because the latter determines what fraction of entrants benefit from the rise of intangibles. For low levels of $G(\bar{A})$ there is little chance that an entrant is highly efficient at intangibles. Because high-intangible firms expand strongly, however, entrants are likely to face a high-intangible incumbent when they attempt to enter. This raises effective entry costs and lowers the incentive to enter. [36] I calibrate $\bar{A}$ by targeting the rise in the average ratio of fixed costs over total costs using the estimates from Section II. For the initial calibration, it is straightforward to show that fixed costs over total costs rise when the homogeneous $\acute{A}$ is calibrated to a lower level. Showing that the relationship between fixed costs and $\bar{A}$ is similarly monotonic is not feasible analytically, as $\bar{A}$ affects the new equilibrium's entry rates, research labor and markup dispersion. A simple comparative static in Online Appendix Figure A10, which plots $\bar{A}$ against average fixed costs over total costs, confirms the existence of a monotonic relationship between fixed costs and $\bar{A}$ in the model's calibration. [37]

In the U.S. calibration, 7.9% of all new firms benefit from a 26% reduction in intangible costs. In the French calibration, 4.4% of all new entrants benefit from a 17% lower $\bar{A}$.[38]

I analyse two experiments on the introduction of these high-intangible firms. In the main experiment, I start with an economy where the share of incumbents with $\bar{A}$ is zero. That is, the rise of high-intangible firms is entirely driven by firms that were not initially operative. This experiment aligns with the observation that the rise of IT-intensity in the 1990s was concentrated in young firms and that the decline of dynamism occurred later for these firms (Haltiwanger, Hathaway and Miranda 2014). In the alternative experiment, I smooth the introduction of high-intangible firms by gradually raising $G(\bar{A})$ from its initial value to its final value over the first 15 years of the transition. I further allow a fraction $G(\bar{A})$ of incumbents in the initial balanced growth path to see the gradual improvement in their intangible costs from the original, homogeneous, level of efficiency to the lower cost $\bar{A}$. The latter assumes that salient differences in intangible costs across firms always existed, but that changes in the availability of technology have made these differences relevant. This experiment aligns with the finding that older firms contributed to the speedup and slowdown in productivity growth since the 1990s in Klenow and

[36] For the U.S., I target the decline in entry in the Business Dynamics Statistics between 1980 and 2016. For France, I impute the decline in entry from the decline in the employment share by entrants in FICUS-FARE, from 1994 to 2016.

[37] To structurally estimate the new parameters, I minimize a loss function similar to (32) with the rise of fixed costs and the decline in entry as moments. To assure that the model does not overshoot the rise in fixed costs, an additional penalty is applied to calibrations where the predicted rise of fixed costs exceeds the empirical rise in fixed costs.

[38] The tables with additional untargeted moments that relate fixed costs to R&D (A10) and markups (A11) show a better match for the final steady state because there is heterogeneity in μ. Fixed effects regressions are still not well-matched, which may be because simulations rely on steady state data. Tables A8 and A9 are well-matched across specifications.

Table 7: Balanced Growth Path Change due to Reduction in Intangible Costs of Top Firms

	Targeted	United States		France	
		¢ Model	¢ Data	¢ Model	¢ Data
Cost Structure					
Average Fixed-Cost Share	Yes	10.6 pp	10.6 pp	4.5 pp	4.5 pp
Slowdown of Productivity Growth					
Productivity Growth Rate	No	-0.3 pp	-0.9 pp	-0.1 pp	-1.3 pp
Aggregate R&D over Value Added	No	34.5%	64.5%	20.1%	5.6%
Decline of Business Dynamism					
Entry rate	Yes	-4.6 pp	-5.8 pp	-1.1 pp	-3.8 pp
Reallocation Rate	No	-36.3%	-23%	-17.6%	-23%
Rise of Market Power					
Average Markup	No	14.6 pt	27 pt	5.6pt	11 pt
Model Wedges					
Labor Wedge	No	6.6 pt	N.A.	3.2 pt	N.A.
Ef ciency Wedge	No	.04 pt	N.A.	.04 pt	N.A.

Notes: Data columns present the empirical moments, while model columns present the theoretical moments. The change in productivity growth is the difference between growth from 1969-1979 (U.S.) or 1969-1994 (France) to growth post-2005. Other U.S. moments equal the difference between 11980 and 2016. Other French moments equal the difference between 1994 and 2016.

Li (2020). Besides their difference in narrative, the main difference between both experiments is that the path of entry is more consistent with the data if high-intangible entrants and incumbents are gradually introduced.

B. Balanced Growth Path

The effect of introducing high-intangible rms is summarized in Table 7, which presents the variables of interest in differences from the original balanced growth path. Two balanced growth path changes are targeted: the increase in xed costs as a percentage of total costs, and the decline in the entry rate. The rise of xed costs is well matched, while the decline in entry is underestimated in both calibrations. The remainder of Table 7 presents results for untargeted objects. These include the slowdown of productivity growth, the decline in business dynamism and the rise of markups. In the U.S. calibration, the model is able to explain about one-half of the rise of markups and one-third of the slowdown of productivity growth. In the French calibration, the model is able to explain most of the decline in the reallocation rate and just over one-half of the rise of markups. The model predicts a 0.1 percentage-point decline in productivity growth. It therefore seems that intangibles are not responsible for most of the slowdown of growth in France. The model overestimates the decline in the reallocation rate in the United States, because all growth in the model occurs through creative destruction. [39]

[39] An empirically relevant additional source of innovation is the improvement of goods that rms already produce (e.g. Garcia-Macia, Hsieh and Klenow 2019, Akcigit and Kerr 2018). In the context of the model, internal innovation would be affected similarly by the rise of intangibles, as high-intangible rms have lower discount rates. In a model like Peters (2020), however, internal innovation primarily raises a rm's market power, hence exacerbating the rise of markups and the decline in wages.

The bottom of Table 7 presents the changes in the labor and efficiency wedges. The labor wedge measures the difference between wages and marginal product, and grew by 6.6 and 3.2 in the U.S. and French calibrations, respectively, due to the rise of markups. The efficiency wedge measures efficiency loss from markup heterogeneity, which increases modestly because low-Á firms have higher markups.

The model predicts a decline in productivity growth despite an increase in aggregate research and development, in line with the data in France and the United States.[40] In a model with homogeneous firms this would be paradoxical, because there is a direct relationship between aggregate R&D and growth. Higher investments and lower growth co-exist in this model because innovation activity is concentrated in a smaller group of high-intangible firms, and because some innovations by low-intangible entrants and incumbents fail to enter the market.

The increase in firm concentration is illustrated in Figure 4, which plots the distribution of firms over the number of products that they produce. This is the most direct measure of concentration in the model. The original balanced growth path is characterized by a lower concentration, featuring more firms that produce one or two goods than is the case in the new balanced growth path. Conversely, the right tail of the firm-size distribution is fatter, indicating that there are more large firms. Note that the increase in concentration is endogenous: high-intangible firms have higher markups and therefore have more incentives to invest in research and development. This causes them to produce a disproportionate fraction of all goods and to grow larger than other firms.

C. Transition Path

The analysis thus far has studied the effect of a rise in intangibles along the balanced growth path. This section shows that short-term dynamics are substantially different. To quantify the transition path, I numerically solve for the path of productivity, markups and wages.[41] This section presents the results from the experiment in which none of the incumbents are assigned a lower intangible cost parameter. The alternative experiment, together with a comparison with the data, are provided in Section E.

The path of productivity growth is presented in Figure 5. Figure 5a presents results for the U.S. calibration, Figure 5b for the French calibration.[42] The solid blue line plots the path of growth in total factor productivity as defined in (24). The yellow dash-dotted line plots the increase in productivity due to the step-wise improvement of quality, which is the source of long-term growth.

When low-Á firms start entering the economy in year 0, there is initially a jump in productivity growth compared to the original steady state (the black upper-dashed line). This is partly because of a rise in entry, driven by the fact that new firms now have a positive probability of being the profitable

[40] The French increase in Table 7 is measured over 1994-2016, while the U.S. increase is over 1980-2016. France experienced a 49.4% increase in R&D over national income between 1980-2016, which is closer to what the model predicts.

[41] The computational algorithm is described in Appendix E.

[42] Figures in the remainder of this section only plot results for the U.S. calibration because the results are qualitatively similar in both calibrations. Full French results are provided in Online Appendix F, Figure A11.

Figure 4. Number of Products before and after a Reduction in Intangible Costs for Top Firms

(a) United States (b) France

Notes: Lines plot the fraction of rms that produce the number of products on the horizontal axis. Solid lines are from the original calibration. Squared lines present the counterpart for the balanced growth path after the introduction of high-intangible rms.

low- Á type, while the high- Á entrants do not face low- Á incumbents yet (Figure 6a). Entry jumps above its initial steady state value as the gures plot the effect of an immediate, permanent shock to G(Á).[43]

As the low-cost rms enter the economy there is a further increase in productivity because they reduce the marginal costs of any good that they produce through the use of intangibles. This causes productivity growth to exceed the growth rate of quality. At peak growth, eight years after the introduction of low- Á entrants, this boosts growth to a level of up to 1.6%. The transitional boom evolves more slowly in France, because a smaller fraction of start-ups bene t from the lower intangible costs (4.4% for France versus 7.9% for the U.S.). The extraordinary growth is predominantly driven by cost reductions from intangibles, consistent with the nding that above-average productivity growth from the mid-1990s to the mid-2000s was primarily caused by IT (Fernald 2015). A slowdown occurs from year 8 onwards in the U.S. calibration. Entry declines because low- Á incumbents produce an increasingly large share of

Figure 5. Transition: Growth Rate of Total Factor Productivity

(a) United States (b) France

Notes: Black- and red dashed lines (respectively) indicate the original and the new steady state.

[43] This is an abstraction, as the true change in the composition of intangible costs was likely more gradual. It may also be the case that some incumbents have seen a reduction in their intangible ef ciency over time. In Section E, I show that entry is in line with the data when entrants and incumbents with low intangible cost parameters are gradually introduced.

Figure 6. Transition Path for Entry, R&D, Markups, Wages

(a) Entry Rate

(b) R&D Intensity

(c) Markups

(d) Wages and Productivity

Notes: Black- and red dashed lines (respectively) in (a) to (c) indicate the original and the new steady state. Figure (a) presents the entry rate, (b) the average ratio of R&D to sales, (c) the average markup, (d) the path of wages (dashed yellow) and productivity (solid blue).

all products. The probability that an entrant benefits from drawing a low Á therefore falls below the probability that it faces a low-Á incumbent, which increases the likelihood of a failed innovation.

The decline in productivity growth is mirrored by an increase in the average ratio of R&D over sales, also known as R&D intensity (Figure 6b). The increase is large: average R&D intensity increases from 2.5 to 7.8%. This is quantitatively very similar to the data. Among U.S. public firms with positive R&D, the average R&D intensity increased from 2.5 (the calibration target) to 8.7%.[44] It aligns with the result that `ideas are getting harder to find' in Bloom et al. (2020), who argue that the effect of innovative investments on growth has diminished. The model offers a potential explanation for their result. As high-intangible firms have higher markups, they have a greater incentive to innovate. Because the returns to R&D are concave, these additional investments have limited effects on growth but increase average R&D intensity considerably, causing the decline in research effectiveness. The presence of low-Á incumbents further means that a fraction of the innovations fail to be introduced to the market, again diminishing the effect of research on growth. Section F elaborates on these inefficiencies.

The model also sheds light on why wages have not kept up with productivity growth in the past 20 years, which has caused a decline in the labor share (Kehrig and Vincent 2021). While the reallocation of economic activity to lower-Á firms leads to a reduction of marginal costs and an increase in produc-

[44] R&D intensity among all public firms increased from 2.0 to 6.7%, again similar to the increase in the model. French R&D expenditure over sales increased from 3.1% among positive spenders (the calibration target) to 4.0%.

Figure 7. Balanced Growth Path Effects of an Increase in Intangible Ef ciency for Top Firms

(a) United States

(b) France

Notes: Balanced growth path growth- and entry rates for various levels of $G(\bar{A})$. Figure 7a plots the U.S. calibration, in which $\bar{A}$ is below the Á of other rms by 26%. Figure 7b plots results for the French calibration, in which $\bar{A}$ is 17% lower than the Á of other rms. Yellow-squared lines present $G(\bar{A})$ in the calibration of Table 7. The lowest $G(\bar{A})$ È 0 plotted is 1%.

tivity, there is no increase in wages because productivity is offset by higher markups (Figure 6c). Note that markups increase because activity reallocates towards high-markup rms, in line with empirical evidence (e.g. Baqaee and Farhi 2020, Autor et al. 2020). This leads to a decoupling of wages and productivity (Figure 6d). Wages continue to grow at the rate of quality improvements, but do not bene t from the transitory increase in productivity growth from intangible adoption.

D. Results: Welfare

To quantify the welfare effect of the rise of intangibles, I calculate the percentage by which consumption along the original balanced growth path must change for the household to be as well off as in the new path of consumption. Two counteracting effects are at play. The initial boom in growth raises the level of productivity, which is positive for welfare. The subsequent slowdown of productivity growth lowers output, which reduces welfare. The permanent rise of R&D worsens this negative effect, because a smaller fraction of the labor force is dedicated to the production of consumption goods.

The overall decline in welfare is sizable: consumption-equivalent welfare falls by 8.92% in the U.S. calibration and by 0.94% in the French calibration (Table 8 columns 1 and 5, row 1). The decline is much larger in the U.S. because steady-state growth falls by more than 0.3 percentage points, versus 0.1 percentage points for France. The temporary increase in growth also lasts longer for France.

The fact that the slowdown of growth is accompanied by a rise in R&D contributes signi cantly to the welfare loss. If L^{rd} is held constant and the new R&D workers assigned back to production, the model predicts a decline in U.S. welfare by 3.2%, close to the welfare loss of 3.3% from the rise in span of control in Aghion et al. (2023). For France, the change in welfare even becomes positive, which means that the initial rise in productivity is suf cient to offset the long-term decline in growth. It follows that the increase in R&D is an important contributor to the welfare effects of the productivity growth slowdown.

The decline in welfare is largely governed by the fraction of rms that have access to the lower intangible costs, $G(\bar{A})$. The effect of $G(\bar{A})$ on growth and entry is illustrated in Figure 7. At $G(\bar{A})$ Æ0, the economy is in the original steady state. As the share of entrants with low-intangible costs becomes positive there is a substantial decline in growth and entry. This is because the smaller $G(\bar{A})$ È 0, the greater the increase in variance and the smaller the reduction of the expected intangible costs. If all rms see a decrease in Á, then average markups would increase, as would the incentive to innovate. A suf ciently homogeneous decrease in Á therefore raises entry and growth above the old steady-state level. Conversely, a mean-preserving spread of Á reduces growth because it reduces incentives to enter. Any change in technology that improves the diffusion of intangibles therefore positively affects entry, growth and welfare. Changes in welfare under alternative calibrations for $G(\bar{A})$ are summarized in the other columns of Table 8. In the main exercise, 7.9% (4.4%) of U.S. (French) rms receive the low intangible cost parameter. The greater the fraction of entrants that receive $\bar{A}$, the higher the welfare. The potential gain is large: if $G(\bar{A})$ is increased to 50%, both calibrations display an increase in welfare of over 20%.

E. Alternative Experiment: Smooth Introduction of Low Intangible Costs

The previous section plotted a transition where the fraction of entrants that receive the low intangible cost parameter $\bar{A}$ jumps overnight, and where all incumbents retain their original intangible costs. I now analyse the transition where, instead, the rise in $G(\bar{A})$ is gradually introduced over 15 years, while the same fraction of incumbents receive the lower $\bar{A}$. The alternative experiment represents, for example, the case in which heterogeneous intangible costs are a salient feature of rms, which becomes relevant when gradual technological advancement enables the use of intangible inputs to reduce marginal costs. This changes the transition path. The alternative experiment has the same calibration for $\bar{A}$ and $G(\bar{A})$ as the main experiment, such that the balanced growth paths are identical.

Results for the U.S. calibration are presented in the left-hand plots of Figure 8, which combine plots from the previous section (solid blue lines) and the alternative experiment (dashed yellow lines). Figure 8a plots productivity growth. When rms with low intangible costs are gradually introduced among entrants and incumbents, the initial increase in growth is slightly smaller. This is primarily due to the behavior of the entry rate, which no longer jumps in the rst year (Figure 8c). Instead, entry increases by a modest 0.7 percentage points over the rst three years of the transition, consistent with occasional increases in entry in the data (Figure 8d). After ten years, the behavior of entry and productivity in both transitions is similar. The rise of markups and R&D intensity are nearly identical in both transitions. Overall, the transition is qualitatively similar with the gradual introduction of low-intangible cost rms. Because the initial boom in productivity is smaller, the alternative transition does raise the fall in welfare: welfare drops by 9.71%, compared to 8.92% in the baseline transition.

Table 8: Welfare Change at Various Levels of Intangible Adoption

	United States				France			
Fraction of low-intangible-cost rms ($G(\bar{A})$):	0.08	0.10	0.25	0.50	0.04	0.10	0.25	0.50
¢ Welfare	-8.92%	-7.46%	2.03%	20.1%	-0.94%	0.49%	8.72%	23.3%
¢ Welfare - constant L^{rd}	-3.18%	-1.48%	8.16%	25.9%	0.54%	2.36%	11.0%	25.5%

Notes: Percentage change from original balanced growth path. $G(\bar{A})$ Æ0.079 for the U.S. and $G(\bar{A})$ Æ0.044 for France in the main analysis.

Figure 8. Transition Path: Model Predictions versus Data (Untargeted)

(a) Productivity Growth (Model)

(b) Productivity Growth (Data)

(c) Entry Rate (Model)

(d) Entry Rate (Data)

(e) Markup (Model)

(f) Markup (Data)

(g) R&D Intensity (Model)

(h) R&D Intensity (Data)

Notes: Black- and red dashed lines (respectively) indicate the original and the new steady state. U.S. calibration. Productivity growth in Figure 8b only includes the smoothed series, as the raw series is highly volatile. HP- lter smoothing parameter is 100. Data sources: productivity growth from Fernald (FRBSF), R&D from Compustat, entry from the BDS, markups from Compustat.

The figures on the right-hand side of Figure 8 plot empirical counterparts to the transition path. The horizontal axes start in 1985 when business dynamism starts trending, and span 45 years to match the theoretical plots. The model is largely able to explain the qualitative features of the data. The model predicts that it takes approximately 45 years for entry rates and markups to converge to the new steady state, while convergence in the data takes 30 years. By that time, the model series have approached levels that are close to their new steady states. The transition duration therefore matches the data.

Productivity growth in the data contains the initial boom and subsequent decline, although its timing differs from the model. The boom occurs between 1995 and 2002, while the model predicts a boom right after high-intangible firms are introduced. The paths of variables in the model also appear more muted than the data, in line with the steady state results. Note, however, that no part of the transition path is targeted. The magnitude of the boom, with growth spiking at 1.6% in the model and 1.7% in the smoothed data, is similar. The boom also lasts for eight years in both series. The model is therefore capable of replicating most qualitative features of the path of productivity growth. Its predictions for entry are closer to the data when a high-intangible firms are gradually introduced, though the paths of R&D and markups are in line with the data in either transition.

Figure A11 in Online Appendix F presents the French transition path, which is qualitatively similar. An empirical comparison is complicated by the fact that data on entry and business dynamism is available only from 1994, so that the effects of intangibles are likely to predate the figures. The ability of the model to fit the time path of growth is worse than for the U.S., furthermore, because productivity growth in France was negative for most years after 2005. The model replicates the duration of the transitions of entry and markups well, though it only explains a modest fraction of the trends for France.

F. Innovation Inefficiencies and Policy

I now take a closer look at the drivers behind the slowdown of productivity growth. The model simultaneously predicts lower productivity growth and higher aggregate R&D expenditures in the final steady state. It is able to do so because the efficiency of R&D declines endogenously. This section describes two new sources of inefficiency that arise when firms have heterogeneous intangible costs, and discusses implications for innovation policy that seeks to remedy them.[45] I first outline these sources and then measure their relative importance.

Sources of Inefficiency

The first inefficiency is that a fraction of innovations are not implemented, because innovators may be undercut by a high-intangible incumbent. I refer to this as the lost innovation channel. The second is that firms of different intangible efficiencies innovate at different rates. Because Ã^{x} È 1, the model features diminishing returns to innovation at the firm level. Differences in innovation rates therefore reduce the average output of researchers. High-intangible firms have higher profits, lower discount rates

[45] The model also features the common sources of innovation inefficiency of models with creative destruction: firms discount the value of their innovations at higher rates than the social planner, causing underinvestment in research. The new channels capture inefficiency in the performance of research, rather than in the aggregate level of research spending.

Table 9: Allocative Efficiency of Researchers: Productivity Growth by Allocation

	United States			France		
	Max.	Actual	Improve (%)	Max.	Actual	Improve (%)
Initial steady state						
1. Keep unsuccessful innovations	1.29	1.29	0.0	1.30	1.30	0.0
2. Reallocate incumbent R&D	1.29	1.29	0.0	1.30	1.30	0.0
3. Reallocate incumbent and entrant R&D	1.29	1.29	0.0	1.30	1.30	0.0
Final steady state						
1. Keep unsuccessful innovations	1.16	0.98	19.4	1.29	1.22	5.7
2. Reallocate incumbent R&D	1.22	0.98	25.5	1.34	1.22	9.8
3. Reallocate incumbent and entrant R&D	1.43	0.98	45.9	1.37	1.22	12.3

Notes: Columns headed `Max.' present maximum productivity growth from reallocating R&D as described in the row. `Actual' presents actual steady-state growth. 'Improve (%)' presents maximum growth as a percentage improvement of actual growth, which measures the efficiency of the actual steady-state allocation of R&D. All rows keep unsuccessful innovation; they assume that the probability that an innovation is successful is 1. Rows `Reallocate incumbent R&D' calculate growth when researchers are reallocated across incumbents, holding the firm-size distribution constant. Rows `Reallocate incumbent and entrant R&D' also reallocate researchers dedicated to entry.

and lower choke prices than low-intangible firms, and therefore invest more in R&D. As the heterogeneous innovation rates are unrelated to a firm's innovation efficiency, low-intangible firms innovate inefficiently little compared to high-intangible firms. I refer to this as research misallocation.

Measuring Inefficiency

I quantify both inefficiency channels by calculating the maximum amount of growth that the number of researchers in the final steady state would generate if the inefficiencies were removed. I hold the firm-size distribution constant at the distribution in the initial and final steady state of the main analysis, such that this section quantifies the direct losses to growth from lost innovation and misallocation.

I conduct three experiments. Experiment 1 quantifies the lost innovation channel. It assumes that each innovation successfully raises the quality of a good, but holds innovation rates across entrants and incumbents (of all sizes) constant. The resulting growth rate is the product of the average innovation step size $\bar{\lambda}$ and the sum of innovation rates from incumbents and entrants. Experiment 2 additionally reallocates researchers in a way that maximizes their innovation output for a given total number of researchers $\bar{L}^{rd}$ that are employed by incumbents. This quantifies the part of the research misallocation channel that is driven by incumbents. I first calculate $\bar{L}^{rd}$, and then allocate them efficiently across firms. The efficient allocation is so that the marginal research output is equal across researchers of differently-sized firms. Note that a firm's intangible efficiency does not affect the number of researchers that it should employ, as research productivity does not depend on a firm's ϕ_i. Experiment 3 extends experiment 2 by additionally reallocating the researchers $\bar{L}^{e}$ that are employed by potential entrants. It thus measures the full extent of the research misallocation channel. The optimal allocations of researchers and resulting growth rates are derived in Online Appendix H.

Figure 9. Entry Rates and Innovation Rates by Firm Size: Actual versus Optimal

(a) United States (b) France

Notes: Figure plots innovation rates $x_n(Á_i)$. Blue-squared lines present the optimal innovation rate according to experiment 3. Red-circled lines present equilibrium innovation rates in the final steady state for low-Á firms, while the yellow-dashed lines present steady-state rates for high-Á firms. The blue triangles are optimal entry according to experiment 3; red triangles are entry in the final steady state.

Results

The growth rate of productivity under each of the three experiments is presented in Table 9. The top panel presents growth rates for the initial calibration in which intangible cost parameters are homogeneous; the bottom panel reports results for the final calibration in which some entrants have higher intangible efficiency. Columns `Max' contain the growth rate $\tilde{g}$ belonging to each of the experiments. Columns `Actual' contain the equilibrium growth rate in that calibration. Columns `Improve (%)' give the percentage by which productivity growth would improve under the experiment's allocation.

The upper panel confirms that the allocation of innovation resources is efficient under homogeneous intangible costs. Firms have equal choke prices, so that all innovations are successful. The marginal research product is also equalized across entrants and incumbents of all sizes, because firms have equal expected profits and creative destruction rates.

In contrast, the bottom panel shows that innovation is inefficient when firms have heterogeneous intangible costs. In particular, experiment 3 in the bottom panel shows that when researchers are efficiently allocated across firms and no innovation is wasted, maximum growth would be 1.43% (1.37%) in the United States' (French) calibration, which is well above growth in the initial steady state. In the actual equilibrium, growth falls from 1.3% to 1% (1.2%) in the United States' (French) calibration, despite a significant increase in $\bar{L}^{rd}$. From comparing the maximum growth in experiments 1 and 3, we learn that both the lost innovation and the research misallocation channel contribute significantly to the overall reduction in innovation efficiency. Experiment 1 shows that if researcher allocations are held constant but the lost innovation is removed, growth would have been 1.16% (1.29%) in the respective calibrations. This means that the misallocation channel comprises 60% (53%) of the total research inefficiency, with the lost-innovation channel accounting for the remainder.

Figure 9 offers a closer inspection of the misallocation drivers. It plots innovation rates against the number of products that firms produce. The figure shows that, compared to the efficient allocation in experiment 3, low-Á (high-intangible) firms innovate excessively. Their high innovation rates are driven by their profitability and low choke price, rather than by a talent for innovation. The figure furthermore shows that the fraction of researchers that is dedicated to entry is insufficient, particularly in the U.S.: the efficient entry rate in experiment 3 is more than twice as high as the actual entry rate in the final steady state. The inefficient entry rate accounts for a large fraction of the misallocation costs to innovation. From comparing experiments 2 and 3 in Table 9, it follows that allowing incumbents' researchers to contribute to entrant innovation yields a much larger improvement in the maximum growth rate than the improvement had from redistributing R&D efficiently across incumbents.

Discussion and Policy Implications

The analysis in this section shows that inefficiency arising from both the lost innovation channel and the misallocation channel causes a significant loss of potential productivity growth. Resolving the inefficiencies is challenging, especially for lost innovation. The firm that produces a good in equilibrium is the firm that, statically, is able to do so at the most favorable terms to consumers. This means that traditional antitrust policy, for example, is unlikely to raise welfare in this setting. The main way to overcome lost innovation is to stimulate the diffusion of intangible technologies. This involves broadening the fraction of entrants that has access to the lower ϕ_i, or creating an efficient market for the innovations that do not make it to production. This, however, falls outside standard policy tools.

Addressing inefficiency arising from the misallocation of researchers may, however, be more straightforward. The misallocation occurs because there is heterogeneity in the private returns to innovation that does not relate to the efficiency with which researchers at the firm produce patents. Policy that alters the private incentives for research can therefore promote a more efficient allocation.It turns out, however, that the subsidy or tax that removes the misallocation inefficiency at a given level of aggregate research effort would have to be sizable. Appendix H analytically derives the rate at which profits would have to be multiplied at entrants and high-ϕ firms to equate their private research incentives. The differences between the required subsidies are substantial. For entrants, the policy would have to multiply the equilibrium private value of innovation by a factor 3.65 (2.09) in the U.S. (French) calibration to bring them in line with the low-ϕ firms. The moderate inequality in intangible efficiency therefore significantly alters the private benefits of research.

V. Extensions

This section explores three extensions. I first show that the model's predictions for productivity growth and business dynamism also hold if markups are constant. I then show that the results in the previous section are robust when firms internalize the diminishing option value of innovation. Finally, I show that the results also hold if intangibles raise fixed costs at the firm level, rather than at the product level.

A. Constant Markups

The analysis thus far has explained the decline in productivity growth and business dynamism jointly with the rise of markups. Recent evidence shows that the labor share in Europe is constant outside of the residential housing sector (Gutiérrez and Piton 2020), while markups may be hard to measure accurately in the absence of data on prices (e.g. Bond et al. 2021).

This section shows that the model predicts a larger decline in productivity growth if markups are constant. To make that point, I impose that all firms charge a constant markup $\bar{\mu}$ over their marginal costs. The markup is calibrated to match the average endogenous markup of 1.29 in the French- and 1.41 in the U.S. calibration of the model. The remainder of the model is left unchanged. In particular, I do not alter the demand system to endogenously arrive at a fixed markup. This facilitates a direct comparison with the main results.The first-order condition for s_{ij} reads $s_{ij} = \left[wY^{-1}\theta\phi_i\bar{\mu}\right]^{\frac{1}{\theta}}$, which follows from inserting the new pricing rule into first-order condition (7). Because markups are homogeneous, the expressions for output and wages respectively simplify to

$$Y = \exp\left(\int_0^1\int 1_{j\in J_i}\ln\left(\frac{q_{ij}}{s_{ij}}\right)di\,dj\right)L^p, \text{ and } w = \exp\left(\int_0^1\int 1_{j\in J_i}\ln\left(\frac{q_{ij}}{s_{ij}}\right)di\,dj\right)\bar{\mu}^{-1}.$$

Online Appendix Table A15 compares the change in the steady-state values in the model with variable markups (columns headed Var. μ) and constant markups (columns headed Fixed $\bar{\mu}$). The rise of high-intangible firms causes productivity growth to fall significantly more when markups are constant: growth now falls by 0.5 and 0.10 percentage points in the U.S. and French calibrations, respectively. When markups are endogenous, high-intangible firms are profitable and invest strongly in R&D. This offsets a part of the decline in growth induced by the fact that high-intangible firms undercut other firms on price. When markups are exogenous there is no motive for R&D by high-intangible firms, worsening the decline in growth. Reallocation rates mirror the additional decline in productivity growth when markups are constant, and now fall well in excess of their empirical decline. In contrast to the data, the table displays a decline in research and development. This is intuitive: constant markups mean that a shift from marginal costs to fixed costs, induced by intangibles, reduces the profit rate. This limits the incentive to expand by high-intangible firms. Combined with the lack of R&D incentives for low-intangible firms due to the low success rate, this explains the R&D decline.

B. Value Function Specification

The preceding analysis relied on a simplified dynamic optimization problem where firms did not internalize the change in their innovation capacity when they added a new product to their portfolio. This assumption significantly improves tractability, as it allows for a closed-form expression of the first-order conditions for innovation. This section shows that the results are qualitatively and quantitatively robust to removing this assumption. The new value function is characterized by

$$rV_t(\phi_i, J_i) - \dot{V}_t(\phi_i, J_i) = \max_{x_i}\left(\begin{array}{c}\sum_{j\in J_i}\pi_t(\phi_i, \lambda_{ij}) + \tau(\phi_i)\left[V_t(\phi_i, J_i\setminus\{\lambda_{ij}\}) - V_t(\phi_i, J_i)\right]\\ + x_i P(\phi_i)\mathbb{E}_{\phi_i}\left[V_t(\phi_i, J_i\cup\lambda^{+}_{ih}) - V_t(\phi_i, J_i)\right] - w_t\eta^x(x_i)^{\psi^x}n_i^{-\sigma}\end{array}\right).$$

The solution of this function is considerably less tractable than the solution in Section I because the function no longer scales linearly in firm size. As firms get larger, the option value of investing in R&D increases, causing them to choose a higher innovation rate. R&D does not fully scale with size, however, because the parameter ¾ is estimated so that the model matches the negative empirical relationship between rm size and growth. Proposition A.1 in Appendix A gives the new value function's solution.

I perform the same experiment as in Section IV. To ease the comparison with the main analysis, I retain most of the previous calibration. I re-estimate ¾ so that the model matches the empirical relationship between rm size and growth. Under an unchanged calibration, the model would predict a strongly negative relationship between rm-growth and rm-size. This is because rms now internalize that the additional option value from producing a good diminishes in n_i. Online Appendix Table 14 details the new model's calibration and main moments. Compared to the original calibration, there is an increase in the value of ¾ for both France and the U.S. The higher parameter value ensures that the empirical deviation from Gibrat's Law is still matched by the model.

Online Appendix Table A16 compares the effect of introducing a calibrated group of high-intangible rms in the model with the new value function speci cation to the effect in the main analysis. The alternative speci cation of the model yields a slightly larger decline in productivity growth of 0.34 percentage points in the U.S. The predicted declines in entry are similar, as are the changes in the reallocation rate. The rise of markups and the changes in the labor and ef ciency wedge are also nearly identical to the main results. These ndings are intuitive - conditional on the recalibration of ¾, the model displays a similar relationship between rm size and rm growth. Because the value function speci cation in this section differs from the value function in the main analysis only in this regard, the results are both qualitatively and quantitatively robust to the use of the full value function.

C. Firm-Level Intangibles

In a nal extension, I explore a model where rms can use their intangibles across each product they produce. Instead of facing a tradeoff between xed- and marginal costs for each product in their portfolio, rms can reduce their marginal costs an equal fraction across all of their products, in exchange for higher xed costs at the level of the rm. This assumption creates strong increasing returns to scale at the rm-level, which are ampli ed when the ef ciency with which rms use intangibles rises.

The model is presented in Online Appendix I. I show that the rise of high-intangible rms causes a qualitatively similar slowdown of productivity growth, fall in business dynamism and rise in markups as the main model does. The effects are quantitatively larger than the results in the main analysis. The model with rm-level intangibles predicts a strong decline of the rate of creative destruction in rm size, so that the model is only solvable if innovation declines rapidly in rm size. I therefore conclude that, while robust to rm-level intangibles, the model in the main text is preferred.

VI. Conclusion

This paper proposes a uni ed explanation for the decline of productivity growth, the fall in business dynamism and the rise of markups. I hypothesize that the rise of intangible inputs — in particular, information technology and software — can explain these trends.

Central to the theory is that intangible inputs shift costs from variable to xed costs, and that rms differ in the ef ciency with which they deploy these inputs. I embed intangibles in an endogenous growth model with heterogeneous multi-product rms, variable markups and realistic entry and exit dynamics. The model suggests that when a subset of new rms becomes more ef cient at using intangible inputs, the aggregate rise of intangibles is accompanied by a decline in entry and long-term growth. I structurally estimate the model to match micro data on U.S. listed rms and the universe of French rms, and nd that intangibles cause a decline of long-term productivity growth of 0.3 percentage points in the U.S. calibration and 0.1 percentage points in the French calibration. Despite the decline of growth, there is an increase in R&D expenditures, in line with empirical evidence. R&D becomes less effective because it is concentrated among a small number of rms and because a fraction of innovators are unable to beat high-intangible incumbents.

While the rise of intangibles negatively affects growth in the long run, its short-run effect is positive. By numerically solving the transition path between the original and the new balanced growth path, I show that growth initially increases for eight years. This is because high-intangible rms initially disrupt sectors by producing goods at lower costs. The overall effect on consumption is negative, although technologies that raise the diffusion of intangibles across rms yield signi cant welfare gains.

References

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. "Innovation, reallocation, and growth." American Economic Review, 108(11): 3450–91.

Adler, Gustavo, Mr Romain A Duval, Davide Furceri, K Sinem, Ksenia Koloskova, Marcos Poplawski-Ribeiro, et al. 2017. Gone with the Headwinds: Global Productivity. International Monetary Fund.

Aghion, P, and P Howitt. 1992. "A Model of Growth through Creative Destruction." Econometrica, 60(2): 323–351.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li. 2023. "A Theory of Falling Growth and Rising Rents." The Review of Economic Studies

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter Klenow, and Li Huiyu. 2022. "Good Rents versus Bad Rents: R&D Misallocation and Growth." Working Paper.

Akcigit, Ufuk, and Sina T Ates. 2021. "Ten facts on declining business dynamism and lessons from endogenous growth theory." American Economic Journal: Macroeconomics, 13(1): 257–98.

Akcigit, Ufuk, and Sina T Ates. 2022. "What happened to US business dynamism?" Forthcoming, Journal of Political Economy.

Akcigit, Ufuk, and William R Kerr. 2018. "Growth through heterogeneous innovations." Journal of Political Economy, 126(4): 1374–1443.

Altomonte, Carlo, Domenico Favoino, Monica Morlacco, Tommaso Sonno, et al. 2021. Markups, intangible capital and heterogeneous nancial frictions. CEP Working Paper.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings. 2019. "International Shocks, Variable Markups, and Domestic Prices." The Review of Economic Studies.

Andrews, Dan, Chiara Criscuolo, Peter N Gal, et al. 2016. "The best versus the rest: the global productivity slowdown, divergence across rms and the role of public policy." OECD.

Atkeson, Andrew, and Ariel Burstein. 2019. "Aggregate implications of innovation policy." Journal of Political Economy , 127(6): 2625–2683.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The fall of the labor share and the rise of superstar rms." The Quarterly Journal of Economics , 135(2): 645–709.

Ayyagari, Meghana, Asli Demirguc-Kunt, and Vojislav Maksimovic. 2018. "Who are America's star rms?" Policy Research working paper no. WPS 8534, World Bank Group.

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson. 2020. "Arti cial intelligence, rm growth, and industry concentration." Firm Growth, and Industry Concentration (November 22, 2020) .

Bajgar, M, C Criscuolo, and J Timmis. 2019. "Supersize Me: Intangibles and Industry Concentration." mimeo .

Baqaee, David Rezza, and Emmanuel Farhi. 2020. "Productivity and misallocation in general equilibrium." The Quarterly Journal of Economics , 135(1): 105–163.

Barkai, Simcha. 2020. "Declining labor and capital shares." The Journal of Finance , 75(5): 2421–2463.

Bartel, Ann, Casey Ichniowski, and Kathryn Shaw. 2007. "How does information technology affect productivity? Plant-level comparisons of product innovation, process improvement, and worker skills." The Quarterly Journal of Economics , 122(4): 1721–1758.

Basu, Susanto. 2019. "Are Price-Cost Markups Rising in the United States? A Discussion of the Evidence." Journal of Economic Perspectives, 33(3): 3–22.

Bessen, James. 2020. "Industry concentration and information technology." The Journal of Law and Economics , 63(3): 531–555.

Bessen, James E, and Cesare Righi. 2019. "Shocking Technology: What Happens When Firms Make Large IT Investments?" Boston Univ. School of Law, Law and Economics Research Paper, , (19-6).

Bijnens, Gert, and Jozef Konings. 2018. "Declining business dynamism in Belgium." Small Business Economics , 1–39.

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb. 2020. "Are ideas getting harder to nd?" American Economic Review, 110(4): 1104–44.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta Eksten, and John Van Reenen. 2014. "IT and Management in America." Working Paper, February 2014.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans do IT better: US multinationals and the productivity miracle." American Economic Review, 102(1): 167–201.

Bloom, Nick, Rachel Grif th, and John Van Reenen. 2002. "Do R&D tax credits work? Evidence from a panel of countries 1979–1997." Journal of Public Economics, 85(1): 1–31.

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch. 2021. "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data." Journal of Monetary Economics.

Bornstein, Gideon. 2018. "Entry and pro ts in an aging economy: The role of consumer inertia." mimeo.

Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt. 2002. "Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence." The Quarterly Journal of Economics, 117(1): 339–376.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. "The productivity J-curve: How intangibles complement general purpose technologies." American Economic Journal: Macroeconomics, 13(1): 333–72.

Bureau of Economic Analysis. 2019. "Data: National Income and Product Accounts (NIPA) series." series Y001RX1A020NBEA, Y006RL1A225NBEA, Y006RX1A020NBEA and GDPCA. Accessed on 2 September 2019 through FRED.

Burstein, Ariel, Vasco Carvalho, and Basile Grassi. 2019. "Data for: Bottom-Up Markup Fluctuations." Working Paper, October 2019. Accessed through CASD project DETVOLT, October 2019 to December 2022.

Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas. 2017. "Rents, Technical Change, and Risk Premia Accounting for Secular Trends in Interest Rates, Returns on Capital, Earning Yields, and Factor Shares." American Economic Review, 107(5): 614–20.

Calligaris, Sara, Chiara Criscuolo, Luca Marcolin, et al. 2018. "Mark-ups in the digital era." OECD Publishing.

Cavenaile, Laurent, Murat Alp Celik, and Xu Tian. 2020. "Are Markups Too High? Competition, Strategic Innovation, and Industry Dynamics." Working Paper.

Corrado, Carol, Charles Hulten, and Daniel Sichel. 2009. "Intangible capital and US economic growth." Review of income and wealth, 55(3): 661–685.

Crespi, Gustavo, Chiara Criscuolo, and Jonathan Haskel. 2007. "Information Technology, Organisational Change and Productivity Growth: Evidence from UK Firms." CEP Discussion Paper No 783, March 2007.

Crouzet, Nicolas, and Janice Eberly. 2019. "Understanding weak capital investment: The role of market concentration and intangibles." *prepared for the Jackson Hole Economic Policy Symposium.*

Davis, Steven J, John Haltiwanger, Ron Jarmin, Javier Miranda, Christopher Foote, and Eva Nagypal. 2006. "Volatility and Dispersion in Business Growth Rates." *NBER macroeconomics annual*, 21: 107–179.

Decker, Ryan A, John Haltiwanger, Ron S Jarmin, and Javier Miranda. 2016. "Where has all the skewness gone? The decline in high-growth (young) firms in the US." *European Economic Review*, 86: 4–23.

Decker, Ryan A, John Haltiwanger, Ron S Jarmin, and Javier Miranda. 2020. "Changing business dynamism and productivity: Shocks versus responsiveness." *American Economic Review*, 110(12): 3952–90.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2014. "The role of entrepreneurship in US job creation and economic dynamism." *Journal of Economic Perspectives*, 28(3): 3–24.

De Loecker, Jan, and Frederic Warzynski. 2012. "Markups and firm-level export status." *The American Economic Review*, 102(6): 2437–2471.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The rise of market power and the macroeconomic implications." *The Quarterly Journal of Economics*, 135(2): 561–644.

De Loecker, Jan, Jan Eeckhout, and Simon Mongey. 2021. "Quantifying market power and business dynamism." *University of Chicago, Becker Friedman Institute for Economics Working Paper*, , (2021-54).

Díez, Federico J, Jiayue Fan, and Carolina Villegas-Sánchez. 2021. "Global declining competition?" *Journal of International Economics*, 103492.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2022. "How costly are markups?" *Forthcoming, Journal of Political Economy.*

Eggertsson, Gauti B, Neil R Mehrotra, and Jacob A Robbins. 2019. "A model of secular stagnation: Theory and quantitative evaluation." *American Economic Journal: Macroeconomics*, 11(1): 1–48.

Farhi, Emmanuel, and François Gourio. 2018. "Accounting for Macro-Finance Trends: Market Power, Intangibles, and Risk Premia." *Brookings Papers on Economic Activity, Fall 2018.*

Fernald, John G. 2015. "Data for: Productivity and Potential Output before, during, and after the Great Recession." *NBER macroeconomics annual*, 29(1): 1–51.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J Klenow. 2019. "How destructive is innovation?" *Econometrica*, 87(5): 1507–1541.

Grossman, Gene M, and Elhanan Helpman. 1991. "Quality ladders in the theory of growth." *The Review of Economic Studies*, 58(1): 43–61.

Gutiérrez, Germán, and Sophie Piton. 2020. "Revisiting the global decline of the (non-housing) labor share." *American Economic Review: Insights*, 2(3): 321–338.

Gutiérrez, Germán, and Thomas Philippon. 2019. "The failure of free entry." NBER.

Hall, Bronwyn H, Jacques Mairesse, and Pierre Mohnen. 2010. "Measuring the returns to r&d." *Handbook of the Economics of Innovation*, 2: 1033–1082.

Haller, Stefanie A, and Iulia Siedschlag. 2011. "Determinants of ICT adoption: Evidence from firm-level data." *Applied Economics*, 43(26): 3775–3788.

Hall, Robert E. 1988. "The relation between price and marginal cost in US industry." *Journal of political Economy*, 96(5): 921–947.

Haltiwanger, John, Ian Hathaway, and Javier Miranda. 2014. "Declining business dynamism in the US high-technology sector."

Harhoff, Dietmar, Frederic M Scherer, and Katrin Vopel. 2003. "Exploring the tail of patented invention value distributions." In *Economics, law and intellectual property*. 279–309. Springer.

Harte-Hanks. 2017. "Ci Technology Database." Access to proprietary this data was kindly provided by Nick Bloom and John van Reenen. Details are provided in the Online Appendix.

Haskel, Jonathan, and Stian Westlake. 2017. *Capitalism without capital: the rise of the intangible economy.* Princeton University Press.

Helpman, Elhanan, and Benjamin Niswonger. 2022. "Dynamics of Markups, Concentration, and Product Span." *American Economic Journal: Macroeconomics*, 14(3): 42–81.

Hopenhayn, Hugo A. 1992. "Entry, exit, and firm dynamics in long run equilibrium." *Econometrica*, 1127–1150.

Hopenhayn, Hugo, Julian Neira, and Rish Singhania. 2022. "From population growth to firm demographics: Implications for concentration, entrepreneurship and the labor share." *Econometrica*, 90(4): 1879–1914.

Hsieh, Chang-Tai, and Esteban Rossi-Hansberg. 2023. "The industrial revolution in services." *Journal of Political Economy Macroeconomics*, 1(1): 000–000.

Hubmer, Joachim, and Pascual Restrepo. 2021. "Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital–Labor Substitution." National Bureau of Economic Research.

Insee. 2019. "Confidential French administrative micro data." The panel of firm-level balance sheet and income statement data derived from FICUS and FARE was kindly provided by Burstein, Carvalho and Grassi (2019). FICUS, FARE, EAE, EAP and TIC data was accessed through CASD project DETVOLT, October 2019 to December 2022. Details are provided in the Online Appendix.

Karabarbounis, Loukas, and Brent Neiman. 2019. "Accounting for factorless income." *NBER Macroeconomics Annual*, 33(1): 167–228.

Karahan, Fatih, Benjamin Pugsley, and Ayşegül Şahin. 2019. "Demographic origins of the startup deficit."

Kehrig, Matthias, and Nicolas Vincent. 2019. "Good dispersion, bad dispersion."

Kehrig, Matthias, and Nicolas Vincent. 2021. "The micro-level anatomy of the labor share decline." *The Quarterly Journal of Economics*, 136(2): 1031–1087.

Klenow, Peter J, and Huiyu Li. 2020. "Innovative Growth Accounting." In *NBER Macroeconomics Annual 2020, volume 35.* University of Chicago Press.

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating firms and aggregate innovation." *Journal of Political Economy*, 112(5): 986–1018.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological innovation, resource allocation, and growth." *The Quarterly Journal of Economics*, 132(2): 665–712.

Koh, Dongya, Raül Santaeulàlia-Llopis, and Yu Zheng. 2020. "Labor share decline and intellectual property products capital." *Econometrica*, 88(6): 2609–2628.

Korinek, Anton, and Ding Xuan Ng. 2019. "Digitization and the Macroeconomics of Superstars." *Working Paper.*

Lashkari, Danial, Arthur Bauer, and Jocelyn Boussard. 2022. "Information technology and returns to scale." *Mimeo.*

Lentz, Rasmus, and Dale T Mortensen. 2008. "An empirical model of growth through product innovation." *Econometrica*, 76(6): 1317–1373.

Levin, Richard C, Alvin K Klevorick, Richard R Nelson, Sidney G Winter, Richard Gilbert, and Zvi Griliches. 1987. "Appropriating the returns from industrial research and development." *Brookings papers on economic activity*, 1987(3): 783–831.

Liu, Ernest, Atif Mian, and Amir Sufi. 2022. "Low interest rates, market power, and productivity growth." *Econometrica*, 90(1): 193–221.

Li, Wendy CY, and Bronwyn H Hall. 2020. "Depreciation of business R&D capital." *Review of Income and Wealth*, 66(1): 161–180.

Mariscal, Asier. 2018. "Firm Organization and Information Technology: Micro and Macro Implications."

Martinez, Joseba. 2019. "Automation, growth and factor shares." *Working Paper, March 2019.*

Martinez, Joseba. 2021. "Putty-Clay Automation." *Working Paper.*

McGrattan, Ellen R. 2020. "Intangible capital and measured productivity." *Review of Economic Dynamics*, 37: S147–S166.

McGrattan, Ellen R, and Edward C Prescott. 2014. "A reassessment of real business cycle theory." *American Economic Review*, 104(5): 177–82.

McKinsey. 2018. "Superstars': The dynamics of firms, sectors, and cities leading the global economy." *McKinsey Global Institute (MGI) Discussion Paper.*

Melitz, Marc J. 2003. "The impact of trade on intra-industry reallocations and aggregate industry productivity." *Econometrica*, 71(6): 1695–1725.

Olmstead-Rumsey, Jane. 2022. "Market Concentration and the Productivity Slowdown."

Peters, Michael. 2020. "Heterogeneous markups, growth, and endogenous misallocation." *Econometrica*, 88(5): 2037–2073.

Peters, Michael, and Conor Walsh. 2021. "Population Growth and Firm Dynamics." *Working Paper.*

Pugsley, Benjamin Wild, and Ayşegül Şahin. 2018. "Grown-up business cycles." *The Review of Financial Studies*, 32(3): 1102–1147.

Rachel, Lukasz. 2021. "Leisure-enhancing technological change." *Working Paper.*

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy*, 98(5).

Salgado, Sergio. 2020. "Technical change and entrepreneurship." *Working Paper.*

Schivardi, Fabiano, and Tom Schmitz. 2019. "The IT Revolution and Southern Europe's Two Lost Decades." *Journal of the European Economic Association.*

Standard, and Poor's. 2019. "Compustat." *Available: Wharton Research Data Services (WRDS)*, Accessed: 30 July 2019.

Syverson, Chad. 2019. "Macroeconomics and Market Power: Context, Implications, and Open Questions." *Journal of Economic Perspectives*, 33(3): 23–43.

Van Reenen, John. 2018. "Increasing differences between firms: market power and the macro-economy." *prepared for the Jackson Hole Economic Policy Symposium.*

Veugelers, Reinhilde, Désirée Rückert, and Christoph Weiss. 2020. "The Growing Digital Divide in Europe and the United States." *European Investment Bank Working Paper.*

Weiss, Joshua. 2020. "Intangible Investment and Market Concentration." *Working Paper.*