Market Concentration and the Productivity Slowdown

Olmstead-Rumsey, Jane

Northwestern University

10 April 2019

Market Concentration and the Productivity Slowdown

Jane Olmstead-Rumsey, Northwestern University

April 10, 2019

Abstract

Since around 2000, U.S. aggregate productivity growth has slowed and product market (sales) concentration has risen. At the same time, productivity differences among firms in the same sector appear to have risen dramatically. In this paper I propose a rich model of competition and innovation to explain the coincidence of these three observations. In the model a key parameter governing all three phenomena is the probability that innovating firms make radical innovations. Thus one explanation for rising concentration, slower productivity growth, and wider technology differences among firms is that the incidence of radical innovations has slowed relative to the 1990s, when the internet and other information technology radically transformed production and sales technology in many sectors.

1 Introduction

Among U.S. public companies, the largest firm's average market share of sales has risen dramatically since the late 1990s (figure 1).[1] Productivity differences between so-called "frontier" or "superstar" firms and their competitors have also been growing since 2000, particularly in information and communications technology (ICT) intensive industries (see figure 2). Despite optimists claiming that the rise of superstar firms will improve allocative efficiency and productivity (Autor et al. (2017b)), aggregate productivity growth has fallen over the same period (also figure 1).[2]

[1] See Grullon et al. (2017) and Council of Economic Advisers (2016) for overviews of trends in market concentration. More than 75% of industries have experienced an increase in the Herfindahl-Hirschman index.

[2] Some argue that productivity growth has not actually slowed, that it has just been persistently mismeasured recently. Syverson (2016) challenges these hypotheses' ability to explain the majority of the measured slowdown using four separate analyses.

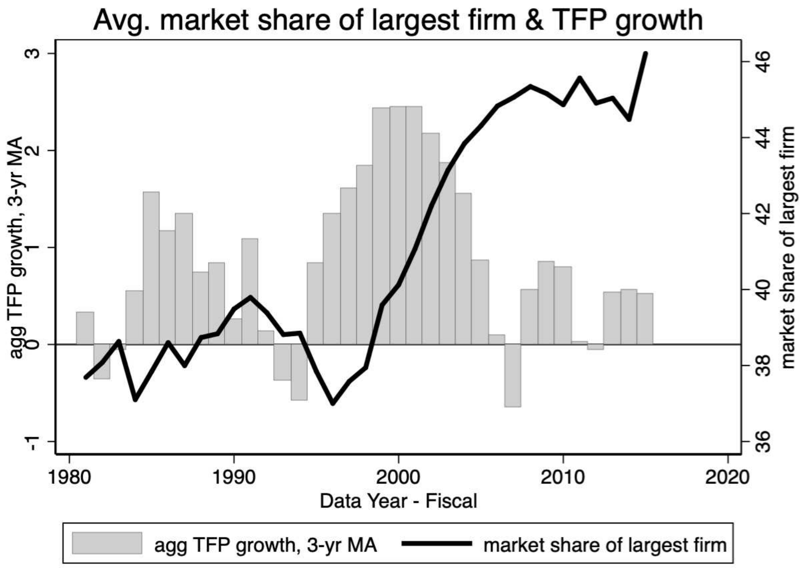

Figure 1: Source: Market share of largest firm (by sales) in 4-digit SIC industries from Compustat (weighted by industry sales); utilization-adjusted total factor productivity (TFP) growth from Fernald (2014).

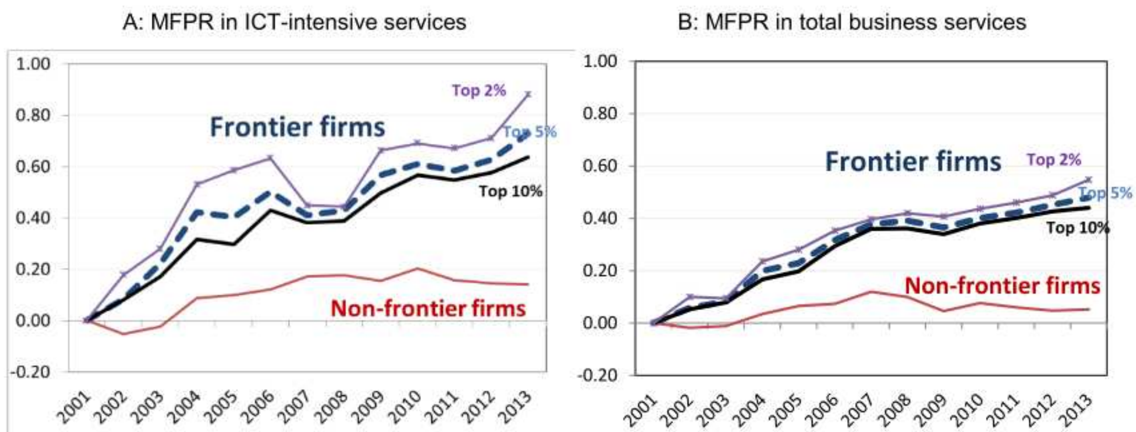

Notes: the global frontier group of firms is defined by the top 5% of companies with the highest MFPQ levels within each 2-digit industry. Laggards capture the all the other firms. Unweighted averages across 2-digit industries are shown for MFPQ, separately for manufacturing and services, normalized to 0 in the starting year. Time period is 2001-2013. Services refer to non-financial business services. ICT-intensive services refer to the information and communication sector (industry code J in NACE Rev. 2) and postal and courier activities (53). See details in Sections 3.3 and 4.4.

Source: Authors' calculations based on the recent update of the OECD-Orbis productivity database (Gal, 2013)

Figure 2: Source: Andrews et al. (2016). MFPR is revenue-based multi-factor productivity. MFPQ is markup-corrected multi-factor productivity.

To explain the coincidence of these three phenomena I propose a general equilibrium quality ladder model of innovation across many sectors. Within each sector, two firms produce products that are imperfect substitutes and interact strategically to set prices and invest in research and development. As in Akcigit et al. (2018), "technology gaps" (quality or productivity[3] differences between firms in the same sector) can take many values, not just neck-and-neck or one-step-ahead/behind as in the seminal work of Aghion et al. (1997). In the model I derive a mapping from quality differences between competitors to market shares, with a firm's market share growing in its *relative* quality. A firm's optimal innovation rate is highest when competitors are neck-and-neck and drops off for both the quality leader and the quality follower as quality differences grow. Aggregate productivity and output growth therefore depend on the distribution of sectors over technology gaps.

In the model, the size of quality improvements conditional on innovating is random. After presenting the model I show that varying the parameter that governs the chance of making a radical innovation (large quality improvement)[4] can match the empirical changes in leader market share, the productivity growth rate, and productivity gaps between firms observed in the U.S. since 2000. Lowering this probability results in endogenously lower innovation effort of firms and greater dispersion of sectors over technology gaps between competitors. Both of these forces contribute to lower aggregate productivity growth, higher leader market share, higher markups, and a lower real interest rate.

Going forward, I plan to use the model presented here to address a variety of questions. First, computing transition dynamics from a steady state with higher probability of radical innovations to lower probability of radical innovations can potentially match the explosion of research and development, productivity, rising margins, and sales growth of large firms in the 1990s and subsequent decline below trend in 2004-2018. Moreover, the model may be helpful in predicting the effects of potentially radical innovations due to artificial intelligence on firm dynamics, market structure, and aggregate productivity growth.

The rest of the paper is organized as follows. In section 2 I discuss the theoretical contribution, the relation to the literature, and present additional empirical motivation. Section 3 presents the model and section 4 lays out preliminary results from a numerical

[3]Higher quality or higher productivity are two sides of the same coin in the model.

[4]The use of "radical innovation" in this paper differs from some other papers in the literature such as Acemoglu & Cao (2015) that use "radical innovation" to refer to an entrant replacing an incumbent.

exercise, comparing the growth rate and other features of the economy in two different steady states with higher and lower probabilities of radical innovations. In section 5 I show that increasing market power in the model is not sufficient to explain the coincidence of the three phenomena that are the focus of the paper.

2 Motivation and Literature Review

2.1 Motivation

Understanding the causes of the productivity slowdown is critical to assessing prospects for future growth and the role that policy can play in alleviating the slowdown. Hall (2015) finds that output in 2013 was 13% below trend (based on 1990-2007) and decomposes this shortfall into various components. Below-trend business investment was the greatest contributor and has been studied by Alexander & Eberly (2018), Crouzet & Eberly (2018), Gutierrez & Philippon (2017), Gutierrez & Philippon (2016), and Jones & Philippon (2016), among others. The second largest contributor was a total factor productivity (TFP) shortfall that accounted for more than a third of the output shortfall and is less well understood.

What theory could connect rising market shares and technological advantages of the top firms with slowing productivity? The "superstar" firm hypothesis and empirical evidence suggest that these large firms are often some of the most productive in their industry, so growth in sales of these firms should increase measured aggregate productivity. But there are also dynamic considerations: when the technology gap between the largest firm and its rivals widens, the large firm might "rest on its laurels" rather than invest in further productivity-enhancing technologies that simply replace its own technology (this is Arrow (1962)'s replacement effect). Thus a dynamic model of productivity growth at the firm level with multiple firms operating heterogeneous technologies in the same sector is needed to untangle the balance of these two forces.

The failure of productive technologies to diffuse to other firms is also a growing concern, according to Anzoategui et al. (2017) and Andrews et al. (2016). Diffusion is an important determinant of productivity growth in firms farther from the technology frontier. With wider technology gaps, smaller firms have a slimmer chance of closing the gap. If the definition of research and development in the model is expanded to include investments with uncertain outcomes, such as attempting to adopt a new technology, the model can also explain this development because it predicts that laggard firms will invest less in quality improvements when quality differences are larger.

2.2 Theoretical Contribution

Formulating a tractable general equilibrium model of strategic interactions in both innovation and pricing decisions is interesting in its own right. Most neo-Schumpterian growth models feature a single firm operating a product line as a monopolist at any given moment in time (see Klette & Kortum (2004), Lentz & Mortensen (2008), Acemoglu & Cao (2015), and Akcigit & Kerr (2018), for leading examples). These models take matching firm-level moments seriously, but are usually silent on industry-level moments. Introducing a duopoly instead allows me to make unified predictions both about market concentration at the industry level and firm-level innovation rates.

This formulation brings together previously distinct strands of literature in macroeconomics concerned with (i) slowing growth (ii) changes in market structure and potentially market power and (iii) superstar firms. Many papers studying the recent rise of large firms have made passing references to the potentially harmful dynamic effects of these large firms on productivity growth but have failed to articulate this link theoretically (OECD (2018)). This model provides a theoretical foundation for the link between the two.

Strands (ii) and (iii) typically rely on opposing assumptions. According to the literature on rising market concentration, incumbent firms exercise greater market power now than in the past and this is reflected in rising markups and profitability (de Loecker & Eeckhout (2017)). On the other hand, the literature on superstar firms typically contends that greater import competition and greater consumer price sensitivity due to better search technology like online retail have *increased* competitive pressures and reduced the market power of incumbent firms, resulting in reallocation to the most productive (superstar) firms (Autor et al. (2017a)). The concept of "market power" is not always well-defined. If measured using markups, as is often done in the literature, the model demonstrates how markups can rise at the same time as there is reallocation to the most productive firms without any changes at all in consumer preferences, including the elasticity of substitution between products.

2.3 Empirical Evidence

The main contribution of the paper is the model, which is presented in section 3. The rest of this section simply provides empirical context. An exploration of the causal relationships among productivity growth, productivity gaps, and concentration is beyond the scope of this paper, but some recent evidence suggests at least a correlation. Gutierrez & Philippon (2017) find that R&D expenditure has slowed down more in more concentrated sectors. Autor et al. (2017b) generally find that sectors with large superstar firms have higher productivity growth over long horizons, but don't investigate changes in the relationship between concentration and productivity growth over time. When taking these changes into account, Gutierrez & Philippon (2017) find positive correlations between concentration and TFP "only before 2002, but an insignificant and sometimes negative correlation after 2002." Autor et al. (2017b) also find that technology diffusion is slower in more concentrated sectors.

2.3.1 Productivity Slowdown

A variety of explanations for the productivity slowdown have been proposed. Most focus on the labor productivity slowdown rather than on total factor productivity. A few explanations are cyclical, such as Anzoategui et al. (2017), who argue that the negative liquidity demand shock that touched off the financial crisis also reduced firms' incentives to introduce new products and adopt existing productive technologies.Such cyclical explanations are unsatisfying for explaining the entire slowdown since the consensus is that the slowdown began well before the global financial crisis.[5]

Secular explanations include the aging workforce (Eggertsson & Mehrotra (2014)) and slowing business dynamism (Decker et al. (2016) and Decker et al. (2018)). Engbom (2017) studies the interactions of aging with innovation and business dynamism. Surprisingly little attention has been devoted to studying firm-level productivity patterns that could illuminate the causes of the productivity slowdown, as I do in this paper.

For example, according to the standard Olley & Pakes (1996) decomposition, aggregate total factor productivity growth could be slowing down for two reasons. First, average TFP growth across all firms could be slowing down. Second, reallocation to the most productive firms (i.e. the growth rate of productive firms) could be slowing down.[6] Appendix A provides details of the firm-level TFP estimation procedure for the

[5]In Anzoategui et al. (2017)'s estimated model of endogenous TFP from 1980 to 2015, they require a negative shock to the productivity of R&D expenditures beginning in the late 1990s to explain why measured R&D in the early 2000s fell below what the model would predict absent the negative R&D efficiency shock. This is similar to the change in the probability of radical innovations I explore in my model.

[6]Formally, as in Olley & Pakes (1996) let a_t be aggregate TFP, $a_{i,t}$ firm-level TFP, and $\bar{a}_t$ unweighted

TFP estimates I refer to throughout the paper. Figure 3[7] shows that the former story is much more important: the fact that the within-firm (unweighted mean component) is what shows a decline means that broad-based below-trend efficiency, not compositional change among U.S. firms, is driving the aggregate slowdown, lending support to explanations focusing on the incentives of existing firms to improve productivity, like the hypothesis I propose here.

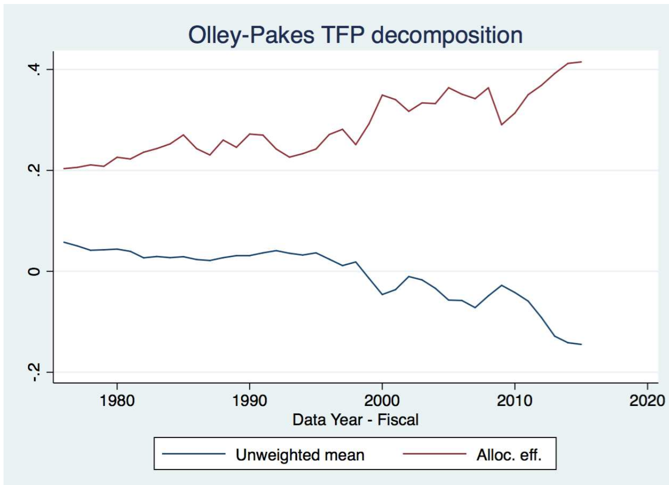

Figure 3: Source: Author's calculations from Compustat.

2.3.2 R&D Slowdown?

The consensus is that innovation and technology adoption drive productivity at the firm, in addition to random shocks (see Griliches (2001) for a survey of the relationship between R&D and productivity at the firm level and Zachariadis (2003) for a leading empirical test). Has the productivity slowdown been accompanied by a productivity

average TFP. Let s denote sales-based market share either individually or on average (with the same notation as for TFP), and N the total number of firms at time t:

$$a_t = \frac{1}{N}\sum_{i=1}^{N} a_{i,t} + \sum_{i=1}^{N}(s_{i,t} - \bar{s}_t)(a_{i,t} - \bar{a}_t)$$

The first term is the unweighted mean and the second is that allocative efficiency term that captures the covariance of size and productivity.

[7] I follow the estimation strategy of de Loecker & Warzynski (2012) to estimate TFP in Compustat, using the variable construction of de Loecker & Eeckhout (2017) in Compustat.

slowdown? Aggregate R&D's share of sales in Compustat has been roughly flat since 1999, before which it had been rising steadily since 1980 (figure 4).

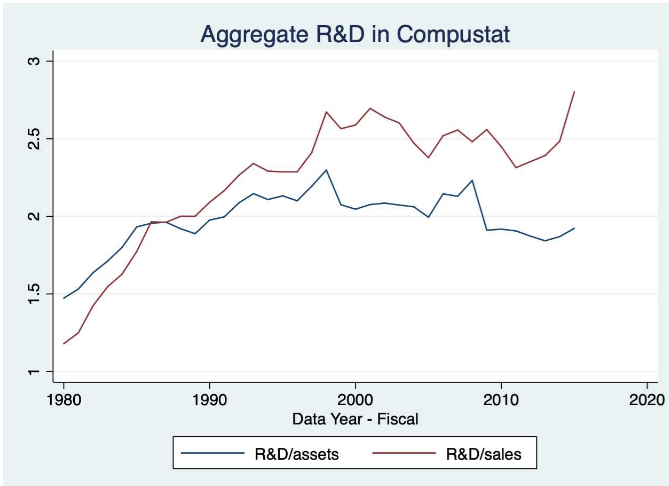

Figure 4: Source: Author's calculations from Compustat summing R&D (XRD) for all firms and dividing by the sum of sales (SALE) or assets (AT) for all firms.

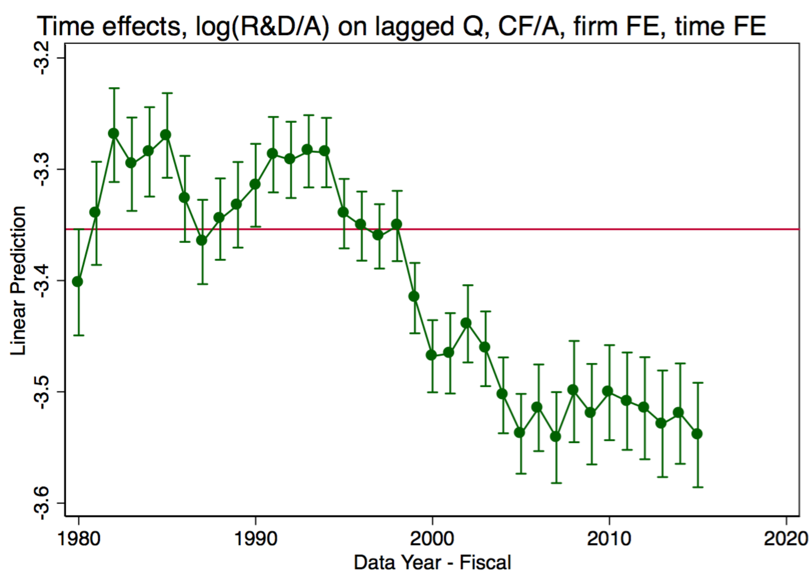

Figure 5: Source: author's calculations from Compustat.

Similar to a firm's decision to invest in physical capital, many factors may influence the decision to invest in innovation. For investment, q-theory tells us that firms should invest when the market value of their assets exceeds the book (replacement) value.[8] I replicate the exercise of Alexander & Eberly (2018) for intangible capital using R&D expenditure as the outcome variable to check whether R&D is slowing down *relative to what theory would predict*. The regression is:

$$\log\left(\frac{R\&D}{assets}_{i,t}\right) = \alpha + \eta_t + X_i + \beta_1 \log\left(\frac{cashflow}{assets}_{i,t}\right) + \beta_2 \log\left(q_{i,t}\right) + \varepsilon_{i,t}$$

I find that R&D has also declined relative to what would be predicted by the theory (figure 5).

In studying the investment slowdown that has partially driven the growth slowdown recently, Gutierrez & Philippon (2016) point out that net investment is low among U.S. public firms despite high value of Tobin's q and explore potential explanations. Increasing concentration appears to be one of the strongest factors correlated with the investment slowdown. They remark that substituting intangible investment or R&D for physical capital investment doesn't change their results. In conclusion, both firm-level average productivity and R&D seem to be slowing down among public firms (with the caveat that q-theory may not be the ideal predictor for optimal R&D behavior).

2.3.3 Productivity Differences

In two companion papers Andrews et al. (2015) and Andrews et al. (2016) summarize characteristics of firms at the global productivity frontier (defined as either the top 100 or the top 5% of firms by estimated productivity in each industry-year). The first paper highlights the growing productivity gap between this frontier and other firms. Globally, frontier firms' productivity has grown at a rate of 3.6% per year while non-frontier firms' productivity grew at just 0.4% over the 2000s.[9] The authors identify two distinct periods: from 2001-2007, frontier firms' productivity grew 4-5% per year and other firms grew 1%, but since the global financial crisis frontier firms have seen productivity growth of just 1% per year while the productivity of the other firms was flat.

To demonstrate that this global fact is also true within the U.S., figure 6 shows average productivity gaps for U.S. public firms in Compustat, either from the industry's market leader in terms of sales or to the most productive firm within its industry.

[8]Formally, I construct Tobin's q in Compustat as $\frac{\text{assets (AT)+shares outstanding (CSHO)*share price (PRCCF)-common equity (CEQ)}}{\text{assets (AT)}}$

[9]Productivity is measured in their papers using both labor productivity and multi-factor productivity to account for the fact that frontier firms tend to be more capital intensive. These numbers refer to labor productivity.

Gaps are defined as the difference between detrended log-TFP of the two firms. I then compute the (unweighted) average of these gaps for all firms within an industry and compute the economy-wide average weighted by industry size. Both measures of the gap are growing over time. The standard deviation of TFP in logs over the sample is 0.8. The average gap to the largest firm grew about 0.4 log points from 1990 to 2010, so about half a standard deviation, and the widening of the gap to the most productive firm was even larger, around three quarters of a standard deviation.

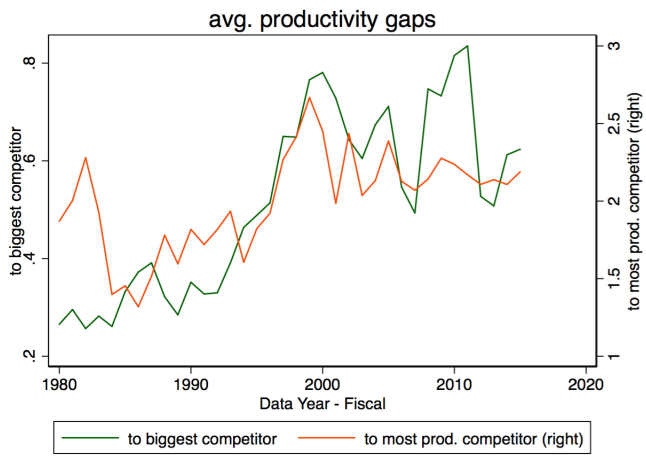

Figure 6: (Unweighted) average productivity gap to either the largest (left axis) or most productive (right axis) competitor in 4-digit SIC, sale-weighted across industries. Compustat.

Averages might not tell the whole story. In figure 7 I show the rightward shift in the distribution of firms over technology gaps to their largest or most productive competitor, verifying that more firms have large technology gaps in 2015 compared to 1995.

2.4 Product Market Concentration

The dramatic rise in the average market share of the largest firm in each industry shown in figure 1 is not just driven by large increases in a few sectors. Figure 8 shows how the entire distribution of industries over the leader's market share has changed from 1995 to 2015. Many more sectors now have just one very large public firm than in 1995, and the peak of the distribution has shifted rightward significantly. This rise is not just due to the increasing market share of foreign firms and resultant mismeasurement

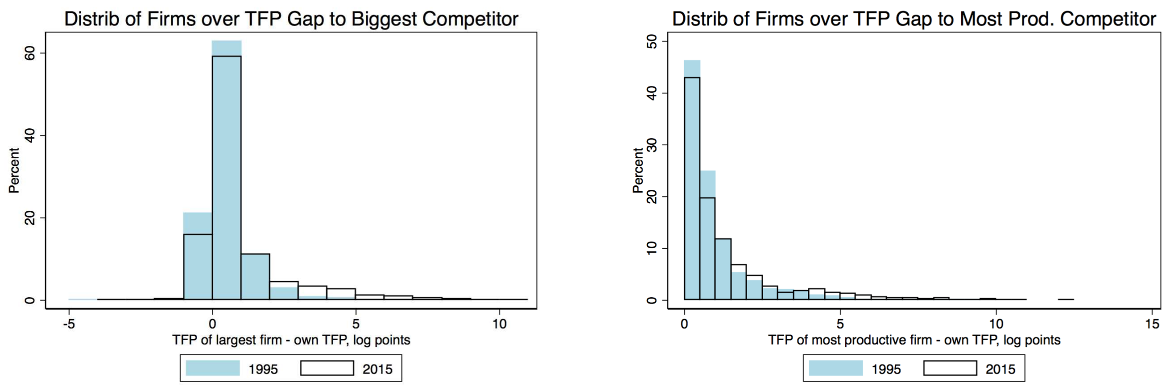

Figure 7: Source: Compustat. Gap is determined to largest (in terms of sales) or most productive firm in 4-digit SIC. A more positive gap means the firm is lagging further behind the leader.

of industry sales in Compustat: Gutierrez & Philippon (2017) constructs an import-adjusted Herfindahl index for the U.S. and a similar rise can be seen in this metric.

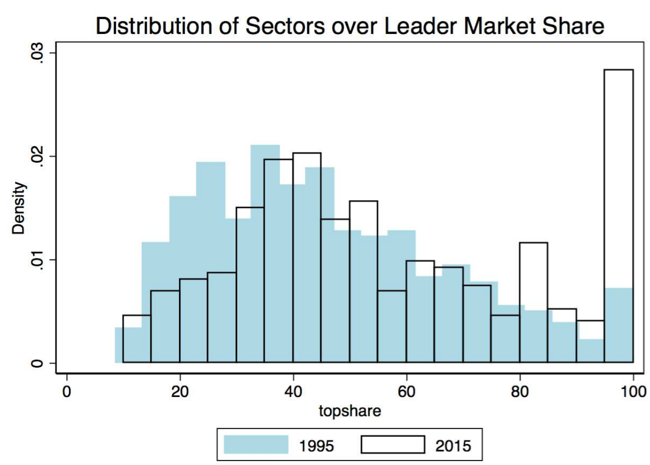

Figure 8: "Sector" refers to 4-digit SIC. Source: Compustat.

A variety of explanations for rising sales concentration have been proposed, from the introduction of ICT that creates winner-take-all markets in a wide variety of industrial classes (retail, entertainment, banking, etc.) and enables the growth of superstar firms (see for example Bessen (2017) and van Reenen (2018)), to excessive regulations that erect barriers to entry and create unnatural monopolies (Gutierrez & Philippon (2017)), to increased mergers and acquisitions (M&A) activity, possibly due to weak antitrust enforcement (Grullon et al. (2017)).

Whatever the cause, rising concentration among the largest firms has been accompanied by widening productivity gaps between these firms and their competitors. More concentrated industries have also seen the greatest slowdown in investment, which has accompanied the aggregate labor productivity slowdown that began in the early 2000s (see Gutierrez & Philippon (2016) and Hall (2015))). Productivity growth has also slowed over this period. To reconcile the coincidence of all three of the phenomena just discussed, the next section presents the model.

3 Model

The model is of a closed economy in continuous time. There are three types agents: a representative household, a representative competitive final good firm, and intermediate goods firms producing capital goods. This section presents the model by going step-by-step through each type of agent in the economy, then analyzes the equilibrium of the model.

3.1 Households

A representative household consumes, saves, and supplies labor inelastically to maximize:

$$U_t = \int_t^{\infty} \exp(-\rho(s-t)) \frac{C_s^{1-\psi}}{1-\psi} ds$$

subject to:

$$r_t A_t + W_t L = P_t C_t + \dot{A}_t$$

I use $\dot{X}$ to denote the time derivative of the variable X.

Households own all the firms, and the total assets in the economy are:

$$A_t = \int_0^1 \sum_{i=1}^{2} V_{ijt} dj$$

Where V_{ijt} is the value of intermediate good firm i in sector j at time t. These value functions are explained in greater detail in section 3.3. The number of firms per sector (two) and the measure of sectors (one) are imposed exogenously. For a balanced growth path with constant growth rate of output g this yields the standard Euler equation $r = g\psi + \rho$.

3.2 Final Goods Producers

The competitive final goods sector combines intermediate goods and labor to create the final output good which is used in consumption, research, and intermediate good production. The final good firm's technology is as follows:

$$Y = \frac{1}{1-\beta}\left(\int_0^1 K_j^{1-\beta} dj\right) L^\beta$$

where K_j is a composite of two intermediate good firms' products within sector j described below. β determines both the elasticity of substitution across sectors ($\frac{1}{\beta}$) and the labor share. For now consider the final good firm's problem of hiring sector composite goods K_j and labor:

$$\max_{K_j, L} P\frac{1}{1-\beta}\left(\int_0^1 K_j^{1-\beta} dj\right) L^\beta - P_j K_j - WL$$

where P is the price of the final good and P_j is the price of the sector j composite good and W is the nominal wage. The first order condition for sector j's composite good given sector j's composite price index P_j will give the demand for sector j's good:

$$K_j = \left(\frac{P_j}{P}\right)^{-\frac{1}{\beta}} L$$

and the real wage is equal to marginal product of labor:

$$\beta\frac{Y}{L} = \frac{W}{P}$$

Now to derive the demand for each firm i within sector j we need to define the sector composite K_j explicitly:

$$K_j = \left((q_1 k_1)^{\frac{\epsilon-1}{\epsilon}} + (q_2 k_2)^{\frac{\epsilon-1}{\epsilon}}\right)^{\frac{\epsilon}{\epsilon-1}}$$

where q_i is the quality of firm i's product (equivalently as firm i's productivity) and k_i is the output of firm i purchased by the final good producer.[10]

Thus within each sector the final good producer will seek to minimize:

$$\min_{\{k_i\}_{i=1}^2} \sum_{i=1}^2 p_i k_i$$

[10]This is the sense in which productivity and quality are equivalent: doubling the quality q_i of both firms has the same effect on final output as doubling the output k_i of both firms.

subject to:

$$K_j = \left((q_1 k_1)^{\frac{\epsilon-1}{\epsilon}} + (q_2 k_2)^{\frac{\epsilon-1}{\epsilon}}\right)^{\frac{\epsilon}{\epsilon-1}} \geq \underline{\mathrm{K}}$$

Taking the first order condition for either k_i yields:

$$p_i = P_j K_j^{\frac{1}{\epsilon}} q_i^{\frac{\epsilon-1}{\epsilon}} k_i^{-\frac{1}{\epsilon}}$$

So, plugging in for K_j we get inverse demand:

$$\begin{aligned} p_i &= P_j \left(\left(\frac{P_j}{P}\right)^{-\frac{1}{\beta}} L\right)^{\frac{1}{\epsilon}} q_i^{\frac{\epsilon-1}{\epsilon}} k_i^{-\frac{1}{\epsilon}} \\ &= k_i^{-\frac{1}{\epsilon}} q_i^{\frac{\epsilon-1}{\epsilon}} P_j \left(\frac{P_j}{P}\right)^{\frac{1}{\beta\epsilon}} L^{\frac{1}{\epsilon}} \end{aligned}$$

Rearranging, the demand function is:

$$k_i = q_i^{\epsilon-1} \left(\frac{p_i}{P_j}\right)^{-\epsilon} \left(\frac{P_j}{P}\right)^{-\frac{1}{\beta}} L \tag{1}$$

3.3 Intermediate Goods Producers

In this section I discuss the production of the intermediate goods and then the innovation technology. In both cases I present the firm's problem first and then the optimal solution. Each intermediate good sector is a duopoly and there is no entry margin. This assumption precludes the possibility of analyzing the contribution of entry to output and productivity growth. Empirical evidence summarized in Bartelsman & Doms (2000) suggests that incumbents are responsible for around 75% of industry-level TFP growth in the U.S. so the model still captures a large share of productivity growth dynamics.

3.3.1 Production and Price Setting

The intermediate goods producers purchase final goods to produce intermediate goods. Each unit of output requires $\eta < 1$ units of the final good to produce. There are no other inputs to intermediate good production.

Facing the demand for their product from the final good producer given in equation 1, I assume the technology *leader* within each sector sets prices a la Bertrand. The *leader* is the firm with higher quality q_i. I assume the *follower* must set price equal to marginal cost η. If the firms are neck-and-neck I assume both set price equal to marginal cost.

This assumption plays an important role in determining the shape of the innovation policy as a function of technology differences, specifically the hump shape. This shape has been suggested theoretically in the work of Aghion et al. (2005), Akcigit et al. (2018), and Schmidt (1997) and found in a variety of studies including Aghion et al. (2005) and Carlin et al. (2004). Intuitively the hump shape appears in this model because the pricing assumption means that the greatest incremental gain in flow profits comes from obtaining quality leadership, so innovation effort will be highest when firms have equal quality (the "neck-and-neck" state).[11] Since this assumption plays an important role in the mechanisms of the model I take a quick detour to discuss it here.

First, the assumption generates empirically plausible predictions about profit shares: the largest U.S. public firms (by sales) capture by far the largest share of industry profits (see figure 9).[12]

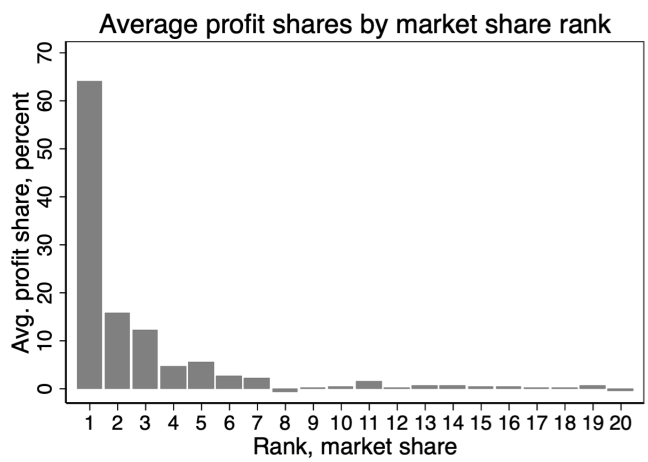

Figure 9: Source: Compustat, 1975-2015. Firms are ranked by market share (sales) within 4-digit SIC industries, and these ranks are compared to profit shares (firm's own operating income as a share of industry-total operating income). The figure averages across 4-digit sectors.

Second, this pricing assumption is actually similar to the assumptions made in other quality ladder models (Klette & Kortum (2004), Acemoglu & Cao (2015), Akcigit & Kerr (2018)), with the caveat that my model features the presence of a lower-quality

[11]Solving the model where both firms set prices a la Bertrand is also possible, but innovation effort is counterfactually low in that model and effort rises smoothly in quality rather than displaying a hump shape.

[12]TFP and size are highly correlated, and the figure looks similar if one uses a productivity ranking instead of sales-based ranks.

substitute to the market leader's product. In those other models, quality improvement over a product line confers a fully enforceable patent on the product until the next innovation occurs. I similarly assume the firm with the higher quality has a fully enforceable patent on its product, but I introduce the notion of sectors and include a second lower quality product in each sector.[13]

If one further assumes that there is no patent protection for the lower quality product and there is free entry to the production of the follower's product, the pricing decisions are the fully optimal outcomes of Bertrand pricing. Under that assumption, the number of followers is indeterminate since the intermediate good production technology is constant returns to scale. The simplest way to resolve this indeterminacy is to assume the lower quality product is produced by a single firm with zero profits.

I now proceed to the solution of the leader's pricing problem. Some definitions are needed. First, sector j's price index:

$$P_j = \left(\sum_{i=1}^{2} q_i^{\epsilon-1} p_i^{1-\epsilon} \right)^{\frac{1}{1-\epsilon}}$$

The market share s_i, plugging in final good firm's demand for k_i:

$$s_i = \frac{p_i k_i}{\sum_{i=1}^{2} p_i k_i} = \frac{q_i^{\epsilon-1} p_i^{1-\epsilon} P_j^{\epsilon-\frac{1}{\beta}} P^{-\frac{1}{\beta}} L}{\sum_{i=1}^{2} q_i^{\epsilon-1} p_i^{1-\epsilon} P_j^{\epsilon-\frac{1}{\beta}} P^{-\frac{1}{\beta}} L} = q_i^{\epsilon-1} \left(\frac{p_i}{P_j} \right)^{1-\epsilon} = \frac{p_i}{P_j} \frac{\partial P_j}{\partial p_i}$$

The final equality holds because:

$$\frac{\partial P_j}{\partial p_i} = \frac{1}{1-\epsilon} P_j^{\epsilon} (1-\epsilon) q_i^{\epsilon-1} p_i^{-\epsilon} = q_i^{\epsilon-1} \left(\frac{p_i}{P_j} \right)^{-\epsilon}$$

Now we are ready to consider the pricing problem of the technology leader i in sector j:

$$\max_{p_i} p_i k_i(p_i) - \eta k_i(p_i)$$

subject to:

$$k_i = q_i^{\epsilon-1} \left(\frac{p_i}{P_j} \right)^{-\epsilon} \left(\frac{P_j}{P} \right)^{-\frac{1}{\beta}} L$$

[13] Another, perhaps more plausible, way to micro-found this assumption is by introducing a cost to filing a patent that is sufficiently high that only the leader, who exercises some additional market power by possessing the higher quality and thus earns higher profits, would be willing to pay. However most Schumpetarian growth models provide no micro-foundation for the monopolist pricing assumption.

taking the first order condition for the price and using the definition of market share above yields the optimal pricing policy:

$$p_i = \frac{\epsilon - (\epsilon - \frac{1}{\beta})s_i}{\epsilon - (\epsilon - \frac{1}{\beta})s_i - 1}\eta$$

So the optimal price is the standard one for nested CES demand: a variable markup that rises in market share. This is easiest to see for the two extreme cases where market share is 0 or 1. When market share is 0, the firm is atomistic with respect to the sector and charges a markup $\frac{\epsilon}{\epsilon-1}$, the CES solution for an elasticity of substitution equal to ϵ. On the other hand, if the market share is 1, the firm only worries about the elasticity of substitution across sectors and sets a markup $\frac{1}{1-\beta} > \frac{\epsilon}{\epsilon-1}$ since products are less substitutable across sectors than within sectors.

It will be important for the (tractable) solution of the model that firms' prices not depend on their quality, only on the *technology gap* between the two firms, defined by the ratio $\frac{q_1}{q_2}$ for firm 1 and $\frac{q_2}{q_1}$ for firm 2. Here I show that this is the case. First, this is clearly satisfied for the technology follower who always sets price equal to marginal cost η regardless of quality.

Second, for the leader, use the definition of the market share and the price index to solve for the market share of the leader i ($-i$ denotes the follower):

$$\begin{aligned} s_i &= q_i^{\epsilon-1}\left(\frac{p_i}{P_j}\right)^{1-\epsilon} \\ &= \frac{q_i^{\epsilon-1}p_i^{1-\epsilon}}{q_i^{\epsilon-1}p_i^{1-\epsilon} + q_{-i}^{\epsilon-1}\eta^{1-\epsilon}} \\ &= \frac{1}{1+\left(\frac{q_{-i}}{q_i}\right)^{\epsilon-1}\left(\frac{p_i}{\eta}\right)^{\epsilon-1}} \end{aligned}$$

Now using the pricing decision of the leader $p_i = \frac{\epsilon-(\epsilon-\theta)s_i}{\epsilon-(\epsilon-\theta)s_i-1}\eta$:

$$s_i = \frac{1}{1+\left(\frac{q_{-i}}{q_i}\right)^{\epsilon-1}\left(\frac{\epsilon-(\epsilon-\theta)s_i}{\epsilon-(\epsilon-\theta)s_i-1}\right)^{\epsilon-1}}$$

Thus there is a mapping from technology gaps to market shares and prices that is independent of quality levels. The market shares and prices for firms with a given technology gap are shown in figure 10. The particular parameterization used to generate the figure is given in table 1 in section 4. The slope of these figures is especially sensitive to ϵ, the elasticity of substitution between firms in the same sector. The leader's optimal price p_i rises slightly as the technology gap widens (that is, as the leader's relatively quality improves). Most of the effect of increased quality appears in the leader's output k_i, so the market share of the leader ($\frac{p_i k_i}{\sum_{i=1}^{2} p_i k_i}$) rises more dramatically in quality than does the price. In this particular parameterization, market share rises from around 32% of sales with one quality step ahead to 38% of the market at 16 steps ahead. The follower, who must sell at $p = \eta$, has a large market share that also increases as the follower's relative quality improves.

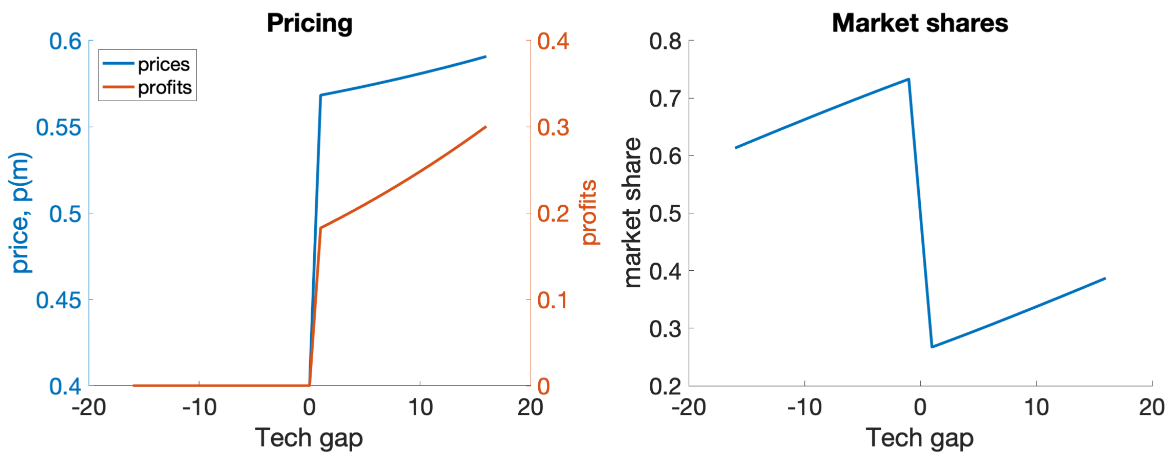

Figure 10: Prices and resulting market shares as a function of the technology gap (ratio of firm qualities).

Obtaining quality leadership in the model generates a drop in market share but, crucially, a rise in profits which is the payoff-relevant object of the firm. Growing market share itself is not an objective of the firm. Of course, in the duopoly setting over some parameter values, including the ones in figure 10, the largest firm in a sales sense is the follower. I focus on the market share of the quality leader in the numerical results because the model's two-firm setup has no direct analogy to industry-level data with more firms per sector. Recall from the discussion of the pricing assumption that an alternative interpretation is that a competitive mass of small firms produce the lower quality product, but a single firm, the quality leader, has the ability to produce the higher quality product.

3.3.2 Innovation

The innovation process for improving the quality of intermediate goods follows Akcigit et al. (2018). Intermediate good firms choose the amount of research spending R of the final good to maximize the discounted sum of all future profits. Innovations arrive randomly at Poisson rate x which depends on research spending according to the function[14]:

$$x = \left(\frac{\gamma R}{\alpha}\right)^{\frac{1}{\gamma}} q_i^{\frac{1-\frac{1}{\beta}}{\gamma}}$$

that is, since $\beta < 1$, at higher quality levels more research spending is needed to achieve the same arrival rate of innovations x. γ and α are R&D technology parameters.

Conditional on innovating the size of the quality improvement is random.[15] Formally, conditional on innovating,

$$q_{i,(t+\Delta t)} = \lambda^{n_{i,t}} q_{i,t}$$

where $\lambda > 1$ is the minimum quality improvement and $n_{i,t} \in \mathbb{N}$ is a random variable. Note that each competitor improves over their own quality and there are no spillovers.[16] Initial qualities of all firms are normalized to 1. Let $N_t = \int_0^t n_s ds$ denote the total number of step size improvements over a product line i since the beginning of time. Define the "technology gap" $m_{1,t}$ between firms 1 and 2 in sector j at moment t as:

$$\frac{q_{1,t}}{q_{2,t}} = \frac{\lambda^{N_{1,t}}}{\lambda^{N_{2,t}}} \equiv \lambda^{m_{1,t}}$$

For numerical tractability I impose a maximal technology advantage $\bar{m}$, but in calibrating the model I will set the parameters so that this maximal gap rarely occurs in steady state.

The probability distribution of possible quality improvements depends on the firm's *relative* quality compared to its competitor. As in Akcigit et al. (2018), I assume there exists a fixed distribution $\mathbb{F}(n) \equiv c_0(n + \bar{m})^{-\phi}$ for all $n \in \{-\bar{m} + 1, \ldots, \bar{m}\}$, shown in the left panel of figure 11, that applies to firms that are the furthest possible distance behind their competitor. The curvature parameter ϕ is critical in the model and determines the speed of catchup by increasing or decreasing the probability of larger innovations. A higher ϕ means a lower probability of these "radical" improvements.[17] c_0 is simply a shifter to ensure $\sum_n \mathbb{F}(n) = 1$.

[14] This specification differs slightly from Akcigit et al. (2018) in the curvature in quality in order to make the firm's solution tractable in the case where both firms produce at once with imperfect substitutes.

[15] Akcigit & Kerr (2018) formalize measures of heterogeneous quality improvements due to innovations using patent citations. They argue that quality improvements are empirically heterogeneous.

[16] Luttmer (2007) provides a rationale for this type of assumption: entrants are typically small and enter far from the productivity frontier, implying that imitation of other firms' technologies is difficult.

[17] As noted by Akcigit et al. (2018), this formulation converges to the less general step-by-step model as $\phi \to \infty$.

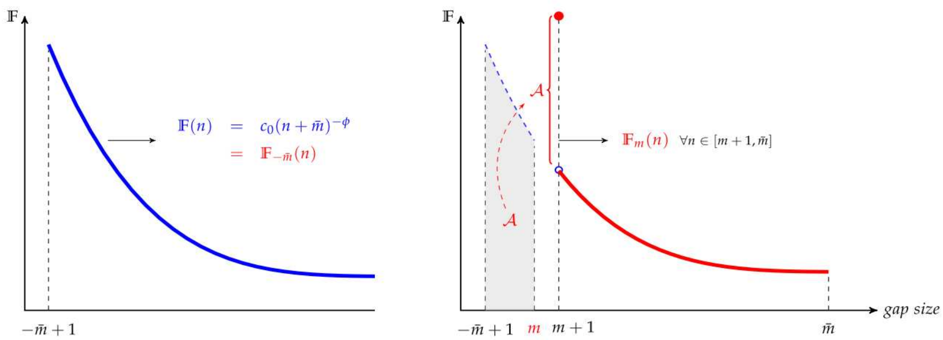

Figure 11: Source: Akcigit et al. (2018)

Given this fixed distribution, the step size distribution specific to each technology gap $m > -\bar{m}$ is given by:

$$\mathbb{F}_m(n) = \begin{cases} \mathbb{F}(m+1) + \mathbb{A}(m) & \text{for } n = m+1 \\ \mathbb{F}(s) & \text{for } n \in \{m+2, \ldots, \bar{m}\} \end{cases}$$

where $\mathbb{A}(m) \equiv \sum_{-\bar{m}+1}^{m} \mathbb{F}(n)$. This distribution is shown in the right panel of figure 11. Simply put, all the mass of the fixed distribution on steps down from current quality are instead put on one-step ahead improvements. The further behind a firm is from its competitor the larger the probability of a more radical improvement.

3.3.3 Value Functions

An intermediate good firm's value function with quality q_t and gap to its rival m_t at moment t is defined as:

$$\begin{aligned} r_t V_{mt}(q_t) - \dot{V}_{mt}(q_t) = \max_{x_{mt}} \{ & \pi(m, q_t) - \alpha \frac{(x_{mt})^{\gamma}}{\gamma} q_t^{\frac{1}{\beta}-1} \\ & + x_{mt} \sum_{n_t = m+1}^{\bar{m}} \mathbb{F}_m(n_t)[V_{nt}(\lambda^{n_t - m} q_t) - V_{mt}(q_t)] \\ & + x_{(-m)t} \sum_{n_t = -m+1}^{\bar{m}} \mathbb{F}_{-m}(n_t)[V_{(-n)t}(q_t) - V_{mt}(q_t)] \} \end{aligned} \tag{2}$$

The firm chooses the arrival rate of innovations x_{mt}. The first line denotes the flow profits and the research cost R_{mt} given the choice of x_{mt}. The second line denotes the probability the firm innovates and sums over the possible states the firm could move to using the distribution of steps sizes and the firm's new value function with higher quality and a larger quality advantage over its rival. The final line denotes the chance the firm's rival innovates and the change in the firm's value because its relative quality falls when the rival innovates.

Now consider the flow profits of the firm, denoting the optimal price of the leader at technology gap m as $p(m)$. We want to eliminate q_t from the value function for tractability, so that each technology gap is associated with a value and the specific firm value function scales in q_t or some function of q_t.

$$\pi(m, q_t) = \begin{cases} 0 & \text{if } m \leq 0 \\ (p(m) - \eta)k_i & \text{for } m \in \{1, \ldots, \bar{m}\} \end{cases}$$

Use the fact that $k_i = q_i^{\epsilon-1} p(m)^{-\epsilon} P_j^{\epsilon - \frac{1}{\beta}} P^{-\frac{1}{\beta}} L$. Normalize P and L to 1. Expanding the definition of the sectoral price index P_j: $k_i = q_i^{\epsilon-1} p(m)^{-\epsilon} (q_i^{\epsilon-1} p(m)^{1-\epsilon} + q_{-i}^{\epsilon-1} \eta^{1-\epsilon})^{\frac{\epsilon - \frac{1}{\beta}}{1-\epsilon}}$. This further simplifies (by dividing by $q_i^{\epsilon-1}$) to:

$$\pi(m, q_t) = \begin{cases} 0 & \text{if } m \leq 0 \\ q_i^{\frac{1}{\beta}-1} (p(m) - \eta) p(m)^{-\epsilon} (p(m)^{1-\epsilon} + (\lambda^{-m})^{\epsilon-1} \eta^{1-\epsilon})^{\frac{\epsilon - \frac{1}{\beta}}{1-\epsilon}} & \text{for } m \in \{1, \ldots, \bar{m}\} \end{cases}$$

So $V_{mt}(q_t) = v_{mt} q_t^{\frac{1}{\beta}-1}$. It can be shown by a guess-and-check approach that this is the case.

The firm's optimal arrival rate x_{mt} (which I refer to as effort in subsequent discussions) is the solution to the first order condition of equation (2), which gives:

$$x_{mt} = \begin{cases} \left(\frac{\sum_{n=m+1}^{\bar{m}} \mathbb{F}_m(n_t)[(\lambda^{n_t - m})^{\frac{1}{\beta}-1} v_{nt} - v_{mt}]}{\alpha} \right)^{\frac{1}{\gamma-1}} & \text{for } m < \bar{m} \\ \left[\frac{1}{\alpha} (\lambda^{\frac{1}{\beta}-1} - 1) v_{\bar{m}t} \right]^{\frac{1}{\gamma-1}} & \text{for } m = \bar{m} \end{cases}$$

The model delivers the predictions that R&D intensity is independent of size (sales) and heterogenous across firms in the same sector, consistent with the empirical evidence discussed in Klette & Kortum (2004).

3.4 Growth Rate of Output

Below I solve for output Y plugging in the intermediate goods firms' output decisions to illustrate the components of output growth:

$$
\begin{aligned}
Y &= \frac{1}{1-\beta}\left(\int_0^1 K_i^{1-\beta} di\right) L^\beta \\
&= \frac{1}{1-\beta}\left(\int_0^1 \left(\sum_{i=1}^{2} q_i^{\frac{\epsilon-1}{\epsilon}} (q_i^{\epsilon-1}\left(\frac{p_i}{P_j}\right)^{-\epsilon}\left(\frac{P_j}{P}\right)^{-\frac{1}{\beta}} L)^{\frac{\epsilon-1}{\epsilon}}\right)^{\frac{\epsilon}{\epsilon-1}1-\beta} dj\right) L^\beta \\
&= \frac{L}{1-\beta} P^{\frac{1-\beta}{\beta}}\left(\int_0^1 P_j^{\epsilon(1-\beta)-\frac{1-\beta}{\beta}}\left(\sum_{i=1}^{2} q_i^{\epsilon-1} p_i^{1-\epsilon}\right)^{\frac{\epsilon(1-\beta)}{\epsilon-1}} dj\right) \\
&= \frac{L}{1-\beta} P^{\frac{1-\beta}{\beta}}\left(\int_0^1 P_j^{-\frac{1-\beta}{\beta}} dj\right)
\end{aligned}
$$

The demand shifter $P^{\frac{1}{\beta}}L$ index is common to all firms and can be taken out entirely (and normalized to one since I assume zero population growth). The price index P_j of each sector falls as the qualities of the two firms in the sector grow, and the exponent is negative for all $\beta \in (0,1)$ so Y is growing in qualities.

Common to all firms in sectors with a particular technology gap m are the prices $p(m)$ of the leader and $p_{-i} = \eta$ of the follower. These prices serve as weights on the qualities of firms within that sector to determine total demand for the sector's products. Thus I can break the problem into groups of firms in sectors with different technology gaps, from $\{0, \bar{m}\}$ (this is a slight abuse of notation as m is the gap perceived by the firm, here I use it to describe the advantage of the leader) and expand the definition of the price index:

$$
\begin{aligned}
Y(t) &= \frac{L}{1-\beta} P^{\frac{1-\beta}{\beta}} \sum_{m=0}^{\bar{m}} \left(\int_0^1 \left(q_{1t}^{\epsilon-1} p_1(m)^{1-\epsilon} + q_{2t}^{\epsilon-1} p_2(m)^{1-\epsilon}\right)^{-\frac{(1-\beta)}{\beta(1-\epsilon)}} \mathbb{1}_{\{j\in\mu_{mt}\}} dj\right) \\
&\equiv \frac{L}{1-\beta} P^{\frac{1-\beta}{\beta}} \sum_{m=0}^{\bar{m}} \left(Q_{mt}\mu_{mt}\right)
\end{aligned}
$$

Here, μ_{mt} is the measure of sectors at each leader's technology gap m at time t and Q_{mt} is a particular weighted average of the qualities of firms within those sectors using the optimal price of a leader with that technology gap and the degree of substitutability across intermediate goods ϵ. The growth rate of output therefore depends on the growth rates $\frac{\dot{Q}_{mt}}{Q_{mt}}$ for each technology gap m which in turn depend on the innovation arrival rates x_{mt} chosen by firms and the exogenous distribution of quality improvement sizes $F(n)$. The final component determining output will be the measure of sectors at each technology gap μ_{mt} that is itself an endogenous object. That is, we can write output:

$$Y = \frac{1}{1-\beta} \sum_{m=0}^{\bar{m}} Q_{mt}(x_{mt}, \mathbb{F}(n)) \cdot \mu_{mt}(x_{mt}, \mathbb{F}(n))$$

3.5 Equilibrium

Let $R_t = \int_0^1 \sum_{i=1}^2 R_{ijt} dj$ denote total research and development spending, C_t total consumption, and $K_t = \int_0^1 \sum_{i=1}^2 \eta k_{ijt} dj$ total purchases of final goods for production of intermediate goods.

An equilibrium in this economy is an allocation $\{k_{ijt}, K_t, x_{ijt}, R_t, Y_t, C_t, L, \mu_{mt}, Q_{mt}\}_{i\in\{1,2\},j\in[0,1],m\in[0,\bar{m}]}^{t\in(0,\infty)}$ and prices $\{r_t, W_t, p_{ijt}\}_{i\in\{1,2\},j\in[0,1]}^{t\in(0,\infty)}$ such that for all t:

1. Intermediate goods firms solve their innovation and price-setting problems (price-setting optimally for the leader only)
2. Final goods firms solve their problem to hire labor and intermediate goods
3. Households solve their consumption-savings problem
4. Goods market clears: $Y_t = C_t + R_t + K_t$
5. Asset market clears, pinning down r_t via the household's Euler equation
6. Labor market clears
7. μ_{mt}, Q_{mt} are consistent with firms' optimal innovation decisions

I focus on a balanced growth path where the measure of sectors with each technology gap μ_{mt} is constant and the growth rate of output is constant. I describe the solution algorithm in the next section.

4 Results

In this section I present preliminary results from a numerical exercise with the model. I briefly describe the numerical solution algorithm for finding the steady state and the choice of the parameters for the numerical exercise before presenting the results. The exercise compares properties of the economy in two different steady states of the model with different values of the probability of radical innovations parameter ϕ. The two values for ϕ are chosen to match average TFP growth rates in 1994-2003 and from 2004-2017. I show that the leader's market share, markups, and productivity gaps are all higher when radical innovations are less likely. The growth rate, average R&D as a share of sales, and the real interest rate are lower.

I then decompose the difference in growth rates in the two steady states into the contributions from the exogenous effect of reducing the average size of quality improvements, the endogenous effect of reduced effort by firms, and the change in the distribution of sectors over technology gaps. I find that these latter two endogenous forces account for about three quarters of the total reduction in the growth rate between the two steady states.

4.1 Solution Algorithm

The algorithm involves first guessing a steady state interest rate. Given this interest rate, solve the value functions for each technology gap by policy function iteration. This process also yields the optimal innovation decisions of leaders and followers at each technology gap. Given these policy functions, initialize quality levels to 1 and iterate forward using the policy functions until the growth rate of the economy stabilizes to a constant value. Finally, check whether this growth rate is consistent with the interest rate guess using the household's Euler equation: $r = g\psi + \rho$. Update the guess of the interest rate and repeat until the interest rate guess and the interest rate implied by the resulting growth rate and the Euler equation are consistent.

4.2 Calibration

For now I take all the parameters of the model from the literature. Eventually I plan to estimate the innovation parameters using a method of moment approach with data on U.S. firms in the two periods of study, using moments like R&D as a share of sales, average TFP differences among firms in the same sector, patenting rates, leader's market share, and the persistence of market leadership at the firm level, plus aggregate growth rates. Table 1 shows the parameterization I use for the current results.

Table 1: Parameters used in the calibration

Parameter	Value	Meaning/source
ψ	2	Intertemporal elasticity of substitution
ρ	0.01	Rate of time preference
β	0.6	Labor share/Nechio & Hobijn (2017)
ϵ	4	Elasticity of substitution within sectors
η	0.4	Marginal cost of intermediate producers
λ	1.49%	Min. qual. improvement, Akcigit et al. (2018)
γ	2	Curvature of R&D fn., Akcigit et al. (2018)
α	0.64	R&D tech., U.S. incumbents, Akcigit et al. (2018)
$\bar{m}$	16	Maximum number of steps ahead

4.3 Numerical Results

I solve for the steady state twice under two different values of the radical innovation parameter ϕ.[18] This is meant to capture the fact that radical innovation probabilities change over time depending on general purpose technologies like the internet and other ICT. Bresnahan & Trajtenberg (1995) identify characteristics of general purpose technologies (GPTs). First, GPTs are *pervasive*, meaning they are applicable in a wide range of sectors. Second, GPTs involve *innovational complementarities*: the productivity of downstream research and development increases as a result of innovation in the GPT. Due to these complementarities, the gains of which are diffuse from the perspective of the sector creating the GPT, their model rationalizes the lags involved in the commercialization of GPTs. For example, the ICT revolution arguably began in the 1970s with the invention of the microprocessor but the biggest gains for productivity growth did not occur until the 1990s. Changing ϕ is an appropriate representation of the changing impact of a GPT since it affects research productivity in all sectors of the economy. The 1990s are often though of as a time of "disruptive" innovations where new innovators made large quality improvements relative to existing products so changing the step size distribution is appropriate.

As a baseline, I use $\phi = 1.5$, which results in a steady state (TFP) growth rate of 1.6% annually, approximately matching average annual TFP growth from 1994-2003. I call this the "quick catchup" regime because under this value of ϕ, firms make larger

[18] I plot the step size distribution for quality improvements under the two values of ϕ in Appendix B.

quality improvements on average conditional on innovating and thus innovating is more likely to result in the laggard catching up and/or overtaking the leader. For the "slower catchup" regime I use $\phi = 1.7$ which results in a steady state annual growth rate of 0.8% to match average TFP growth frmo 2004-2017.

Table 2: Steady state comparison

Moment	$\phi = 1.5$	$\phi = 1.7$
Annual growth rate	1.6%	0.78%
Ann. ret. on assets, r	4.2%	2.5%
Avg. leader market share	30.6%	31.5%
Avg. tech. gap, m	5.6	7.0
Avg. R&D/sales	4.0%	2.0%
Avg. markup	1.43	1.44

Table 2 compares features of the equilibrium in the two steady states. A rise in the average productivity/technology gap drives a modest rise in leader market share and markups[19] in the slower catchup steady state. For market concentration, the 1 percentage point rise in leader market share is about 12.5% of the total rise from 1998 to 2017. The magnitude of changes in leader market share and markups are highly sensitive to the choice of substitutability ϵ between goods in the same sector, and grow as ϵ grows. The lower probability of radical innovation induces lower innovation effort by firms: average R&D as a share of sales falls from 4% to 2%.[20] Finally, consistent with the lower growth rate, r, the risk-free rate of return/rate of return on the portfolio of firms, falls from 4.2% annually to 2.5%.

In figure 12 I plot the firm innovation policy functions under the two regimes as a function of the technology gap to the firm's competitor. The main difference between the two steady states is the innovation effort of the follower, which is slightly lower in the slow catchup regime. This can also be seen as corresponding to slower technology adoption and diffusion of laggard firms observed in the data. This change in effort, combined with the exogenous change in ϕ, results in the stationary distribution of firms

[19] The average markup for U.S. public firms estimated in de Loecker & Eeckhout (2017) went from around 1.4 in the late 1990s to 1.67 in 2014, so the model explains only a small share of the total change.

[20] Consistent with evidence summarized in Klette & Kortum (2004), R&D in this model scales perfectly in firm size (measured by sales), as long as the technology gap to the competitor is held fixed.

over technology gaps shown in the second panel of the figure. Many fewer sectors are near the neck-and-neck state at any given time in the slower catchup regime.

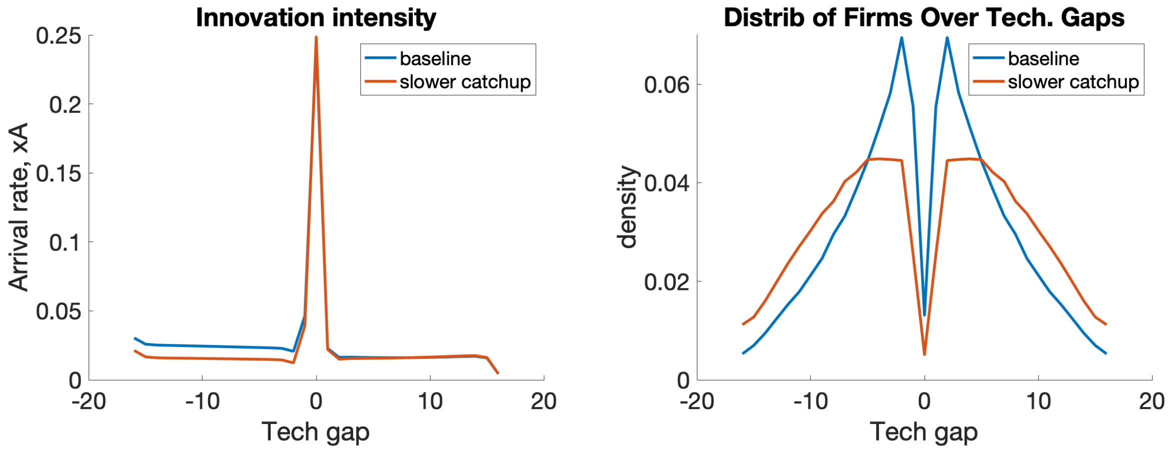

Figure 12: Firm policy functions depending on technology gap (a) and stationary distribution of firms over technology gaps (b), two catchup regimes.

4.4 Decomposition of Growth Rate Change

Recall that the growth rate depends on several factors: the probability of radical innovations ϕ which governs the size of quality improvements conditional on innovating; the endogenous innovation effort of firms x_{mt}, which governs the arrival rate of innovations; and the endogenous distribution of firms over technology gaps μ_{mt}, which determines the size of the mass of firms choosing each innovation rate.

$$Y = \frac{1}{1-\beta} \sum_{m=0}^{\bar{m}} Q_{mt}(x_{mt}, \mathbb{F}^{\phi}(n)) \cdot \mu_{mt}(x_{mt}, \mathbb{F}^{\phi}(n))$$

Table 3 decomposes the contribution of each of these components to the change in growth rates between the two steady states. Keeping innovation effort and the distribution of firms over technology gaps fixed, there is a first-order effect on growth from reducing the probability of radical innovations. This lowers the growth rate from 1.6% to about 1.4% annually. The much larger share of the effect comes from the mechanisms of the model. First, firms respond to the lower probability of large quality improvements by reducing innovation effort. About half the decline in the growth rate is due to this effect. Second, the first two factors make the distribution of firms over technology gaps more spread out. Since innovation effort of both the leader and the follower is lower at wider technology gaps, this contributes a further 18 basis points to the decline in the growth rate.

Table 3: Decomposition of growth rate change

	g	Contrib.
Baseline	1.60%	
+ step size distrib., $\mathbb{F}^{\phi}(n)$	1.38%	27%
+ innov. effort, x_{mt}	0.96%	51%
+ firm distrib., μ_{mt}	0.78%	22%

4.5 Discussion

One reason to believe this exercise comparing two steady states with different probabilities of radical innovations is the correct one to fit the situation in the U.S. since the 1990s is simply that the results are consistent with what I observe empirically. Targeting just the growth rate delivers the correct direction of the non-targeted moments including leader market share and productivity gaps. Under the slow catchup regime, the stationary distribution of firms has wider tails and less mass around the neck and neck position. This generates the model's analogy to figures 7 and 8, where more sectors have larger technology gaps and higher leader market share. Recall also that figure 2 shows that productivity divergence is particularly pronounced in ICT intensive sectors, pointing to ICT as a potential cause. Bessen (2017) also shows that ICT intensity is correlated with rising market concentration.

A second reason is that empirical evidence shows the "advantage of backwardness" is lessening over time, consistent with the idea that it is harder to make large productivity improvements through innovation. Andrews et al. (2016) show that in a regression of firm-level productivity growth on a variety of explanatory variables, the coefficient on the lagged productivity gap to the technology frontier has been declining over the 2000s (see figure 13), potentially consistent with a change in the step size distribution in the model.[21]

[21] However, this empirical observation is endogenous according to the model, because it may be a result of both structural change to catchup speeds and to the endogenously lower innovation effort by laggard firms since their regression doesn't control for innovation effort (R&D investment).

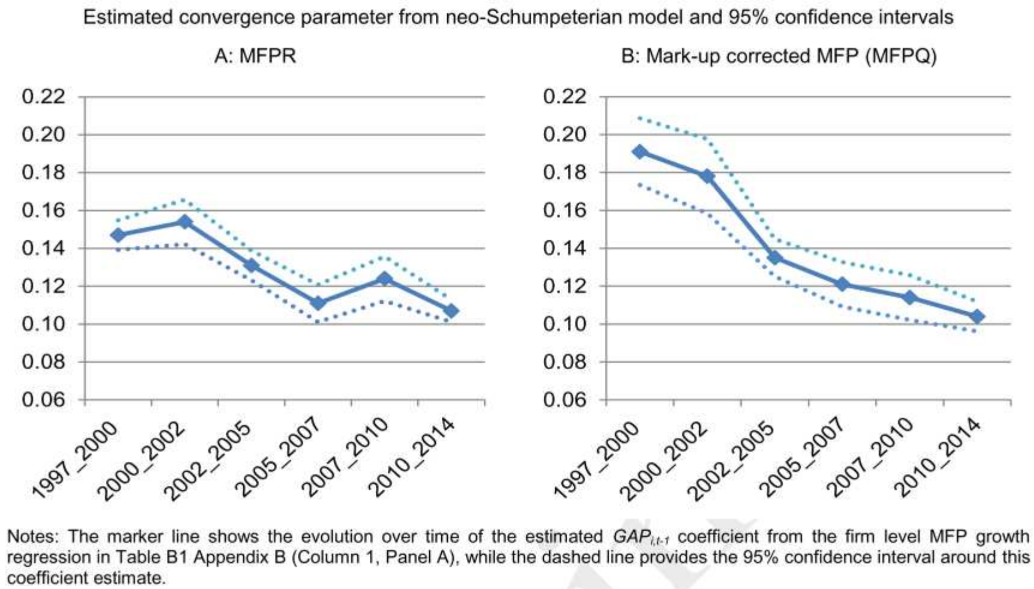

Figure 13: Speed of technological convergence is slowing down. Plotted coefficient is the coefficient on a firm's lagged multi-factor productivity (MFP) gap in a regression of current MFP on lagged MFP gap to the frontier. Source: Andrews et al. (2016).

5 Other Results

5.1 Increasing Market Power?

Recent research has focused on the potential costs of rising market power and markups (see de Loecker & Eeckhout (2017), Eggertsson et al. (2018) and Edmond et al. (2018) for example) for growth and welfare. Can an increase in market power generate the same predictions for the macroeconomic changes experienced in the U.S. in recent years as a change in the probability of radical innovations? In this section I show that in this model it cannot. I model an increase in market power as a decrease in the substitutability of products in the same sector, ϵ.

I keep the calibration the same as in table 1 and set $\phi = 1.5$ (1990s case). I decrease ϵ from 4 in the baseline to 3 to approximately match the current markup among U.S. public firms estimated by de Loecker & Eeckhout (2017). Because the market leader faces a lower elasticity of demand and charges a higher price when market power increases, the average leader market share of sales rises slightly compared to the baseline despite the fact that average quality differences (technology gaps) between competitors fall. Other than this, the results are the opposite of what we observe in the data. Because of greater market power, profits are higher when the firm has market leadership and this induces higher innovation effort at all levels of the technology gap (average R&D/sales rises from 4.6% to 7.8%, see table 4 and figure 14). This results in a higher growth rate and higher interest rate. There is greater turnover in market leadership and average productivity differences go down (figure 14).

Table 4: Steady state comparison, market power

Moment	$\epsilon = 4$	$\epsilon = 3$
Annual growth rate	1.6%	2.3%
Ann. ret. on assets, r	4.2%	5.6%
Avg. leader market share	30.6%	31.0%
Avg. tech. gap, m	5.6	5.1
Avg. R&D/sales	4.0%	7.8%
Avg. markup	1.43	1.62

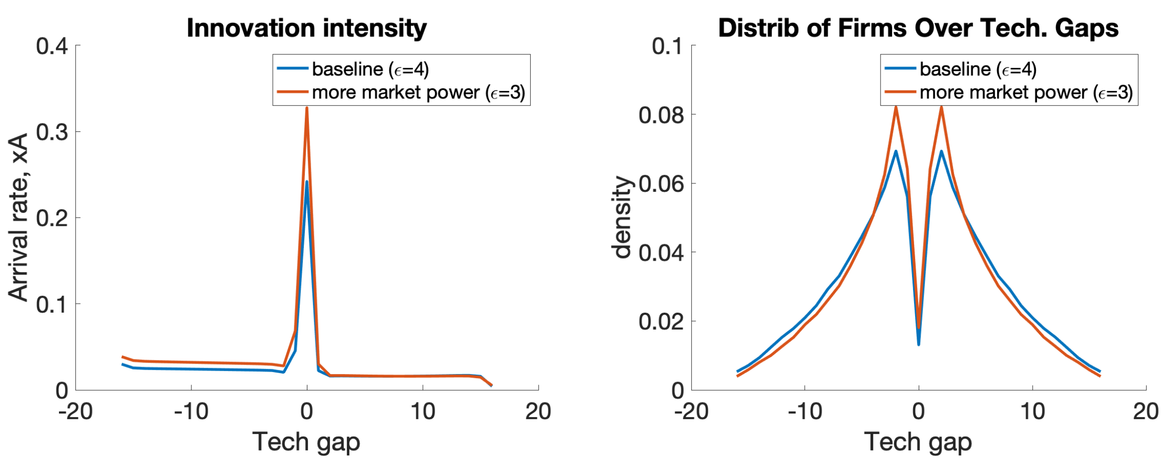

Figure 14: Firm policy functions depending on technology gap (a) and stationary distribution of firms over technology gaps (b), two market power regimes.

6 Conclusion

I have presented a general equilibrium model of innovation and growth where multiple firms are active in a sector in each period and goods within sectors are imperfect substitutes. Future work will involve the estimation of the model for U.S. firm-level data and computation of transition dynamics of the economy transitioning from high to low radical innovation steady states.

The model is able to match changes in the distribution of sectors over technology gaps and leader market shares from the 1990s to the 2010s with a change in the probability of radical innovations. An increase in market power cannot generate the same fit for data. The model jointly explains rising concentration, increasing productivity differences between firms in the same sector, and the productivity slowdown. It can account for what might be called the "superstar productivity puzzle": despite the rapid growth of a few highly productive and already large firms over the 2000s, *aggregate* productivity growth has slowed down rather than sped up. The model predicts that this is because of the dynamic effects of a dominant leader on the innovation decisions of laggard competitors.

Through the lens of the model I unify the Schumpeterian endogenous growth literature with the growing literature on rising concentration, markups and market power in the U.S. The model demonstrates a novel effect of general purpose technologies: rising market concentration that can be misinterpreted as rising market power if not accounting for changes in productivity differences between firms.

References

Acemoglu, D., & Cao, D. (2015). Innovation by entrants and incumbents. *Journal of Economic Theory*, *157*, 255–294.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., & Howitt, P. (2005). Competition and Innovation: An Inverted-U Relationship*. *Quarterly Journal of Economics*, *120*(2), 701–728.

Aghion, P., Harris, C., & Vickers, J. (1997). Competition and growth with step-by-step innovation: An example. *European Economic Review*, *41*(3-5), 771–782.

Akcigit, U., Ates, S., & Impullitti, G. (2018). Innovation and Trade Policy in a Globalized World. *NBER Working Paper*, (24543).

Akcigit, U., & Kerr, W. R. (2018). Growth Through Heterogeneous Innovations. *Journal of Political Economy*, *126*(4).

Alexander, L., & Eberly, J. (2018). Investment Hollowing Out. *IMF Economic Review*, *66*(1), 5–30.

Andrews, D., Criscuolo, C., & Gal, P. (2016). The Global Productivity Slowdown, Technology Divergence and Public Policy: A Firm Level Perspective. *The Future of Productivity: Main Background Papers*, (pp. 1–50).

Andrews, D., Criscuolo, C., & Gal, P. N. (2015). Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence From OECD Countries. *The Future of Productivity: Main Background Papers*, (p. 38).

Anzoategui, D., Comin, D., Gertler, M., & Martinez, J. (2017). Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence. *NBER Working Paper*, (22005), 1–52.

Arrow, K. J. (1962). The Economic Implications of Learning by Doing. *The Review of Economic Studies*, *29*(3), 155–173.
URL `http://www.jstor.org.ezproxy.unal.edu.co/stable/pdf/2295952.pdf`

Autor, D., Dorn, D., Katz, L. F., & Patterson, C. (2017a). Concentrating on the Fall of the Labor Share. *107*(1476), 180–185.

Autor, D., Katz, L. F., & Dorn, D. (2017b). The Fall of the Labor Share and the Rise of Superstar Firms. *NBER Working Paper Series*, (23396).

Bartelsman, E. J., & Doms, M. E. (2000). Understanding Productivity: Lessons from Longitudinal Microdata. *Journal of Economic Literature*, *XXXVIII*(September), 569–594.

Bessen, J. (2017). Information Technology and Industry Concentration. *Boston University School of Law Law & Economics Paper*, *17*(41).

Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies 'Engines of growth'? *Journal of Econometrics*, *65*(1), 83–108.

Carlin, W., Schaffer, M., & Seabright, P. (2004). A minimum of rivalry: Evidence from transition economies on the importance of competition for innovation and growth. *Contributions to Economic Analysis and Policy*, *3*(1), 241–285.

Council of Economic Advisers (2016). Benefits of Competition and Indicators. *Issue Brief*, (May).

Crouzet, B. N., & Eberly, J. (2018). Intangibles, Investment, and Efficiency. *AEA Papers and Proceedings*, *108*, 426–431.

de Loecker, J., & Eeckhout, J. (2017). The Rise of Market Power and the Macroeconomic Implications. *Working Paper*, (pp. 1–44).

de Loecker, J., & Warzynski, F. (2012). Markups and Firm-Level Export Status: Appendix. *American Economic Review*, *102*(6), 2437–2471.

Decker, R. A., Haltiwanger, J., Jarmin, R. S., & Miranda, J. (2016). Where has all the skewness gone? The decline in high-growth (young) firms in the U.S. *European Economic Review*, *86*, 4–23.

Decker, R. A., Haltiwanger, J., Jarmin, R. S., & Miranda, J. (2018). Changing Business Dynamism and Productivity: Shocks vs. Responsiveness.

Edmond, C., Midrigan, V., & Xu, D. Y. (2018). How Costly Are Markups? (September).

Eggertsson, G., & Mehrotra, N. (2014). A Model of Secular Stagnation. *NBER Working Paper*, (20754).

Eggertsson, G., Robbins, J., & Wold, E. G. (2018). Kaldor and Piketty's Facts: The Rise of Monopoly Power in the United States. (pp. 1–69).

Engbom, N. (2017). Firm and Worker Dynamics in an Aging Labor Market. *Working Paper*, (pp. 1–115).

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. *Federal Reserve Bank of San Francisco*.

Griliches, Z. (2001). *R&D, Education, and Productivity: A Retrospective*. Cambridge, Mass.: Harvard University Press.

Grullon, G., Larkin, Y., & Michaely, R. (2017). Are US Industries Becoming More Concentrated? *Working Paper*, (pp. 1–68).

Gutierrez, G., & Philippon, T. (2016). Investment-Less Growth: an Empirical Investigation. *NBER Working Paper*, (22897).

Gutierrez, G., & Philippon, T. (2017). Declining Competition and Investment in the U.S. *NBER Working Paper*, (23583).

Hall, R. E. (2015). Quantifying the Lasting Harm to the US Economy from the Financial Crisis. *NBER Macroeconomics Annual*, *29*(1), 71–128.

Jones, C., & Philippon, T. (2016). The Secular Stagnation of Investment? *Working Paper*, (December), 1–25.

Klette, T. J., & Kortum, S. (2004). Innovating Firms and Aggregate Innovation. *Journal of Political Economy*, *112*(5), 986–1018.

Lentz, R., & Mortensen, D. (2008). An Empircal Model of Growth Through Product Innovation. *Econometrica*, *76*(6), 1317–1373.

Luttmer, E. G. (2007). Selection, Growth, and the Size Distribution of Firms. *Quarterly Journal of Economics*, (August), 1103–1145.

OECD (2018). Market Concentration. *OECD Issue Paper*, *46*, 1527–1527.

Olley, G. S., & Pakes, A. (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, *64*(6), 1263–1297.

Schmidt, K. M. (1997). Managerial Incentives and Product Market Competition. *The Review of Economic Studies*, *64*(2), 191–213.

Syverson, C. (2016). Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown. *Working Paper*, (February), 1–28.

van Reenen, J. (2018). Increasing Differences Between Firms: Market Power and the Macro-Economy. *CEP Discussion Paper*, (1576).

Zachariadis, M. (2003). R&D, innovation, and technological progress: a test of the Schumpeterian framework without scale effects. *Canadian Journal of Economics/Revue Canadienne d'Economique*, *36*(3), 566–586.

A TFP Estimation

I use Compustat data on U.S. public firms from 1975-2015 to estimate total factor productivity (TFP) at the firm level. I focus on the non-financial sector and exclude utilities and firms without an industry classification. I keep only those companies that are incorporated in the U.S. The sample includes around 3,000 firms per year.

I construct each firm's capital stock $K_{i,t}$ by initializing the capital stock as PPEGT (total gross property, plant, and equipment) for the first year the firm appears. I then construct $K_{i,t+1}$ recursively:

$$K_{i,t+1} = K_{i,t} + I_{i,t+1} - \delta K_{i,t}$$

where PPENT (total net property, plant, and equipment) is used to capture the last two terms (net investment). I deflate the nominal capital stock using the Bureau of Economic Analysis (BEA) deflator for non-residential fixed investment.

In de Loecker & Warzynski (2012) the authors show that under a variety of pricing models the firm's markup can be computed as a function of the output elasticity θ_{it}^V of the variable input and the variable input's cost share of revenue[22] :

$$\mu_{it} = \theta_{it}^V \frac{P_{it} Q_{it}}{P_{it}^V V_{it}}$$

Following de Loecker & Eeckhout (2017) I use COGS (cost of goods sold) deflated by the BEA's GDP deflator series as the real variable input cost $M_{i,t}$ of the firm. While the number of employees is well measured in Compustat and would be sufficient to estimate productivity, the wage bill is usually not available and would be needed to compute the labor cost share needed to compute the markup simultaneously with productivity. I don't use this markup information in the current draft but plan to use it to differentiate between TFPR (revenue-based TFP) and TFPQ (markup-corrected TFP).

For the results presented in this paper, I assume a Cobb-Douglas production function[23] for firm i in 2-digit SIC sector s in year t so that factor shares may vary across sectors:

[22]This approach requires several assumptions. First, the production technology must be continuous and twice differentiable in its arguments. Second, firms must minimize costs. Third, prices are set period by period. Fourth, the variable input has no adjustment costs. No particular form of competition among firms must be assumed.

[23]Below I show the correlation of these firm-level TFP estimates with estimates from a model assuming a translog production function.

$$Y_{i,s,t} = A_{i,s,t} M_{i,s,t}^{\beta_{M,s}} K_{i,s,t}^{\beta_{K,s}}$$

I use the variable SALE to measure firm output $Y_{i,s,t}$. I deflate SALE using the GDP deflator series to obtain real output at the firm level. I use time dummies η_t to detrend the TFP estimates and obtain TFP in logs (lower case variables denote variables in logs) by computing the residual of the following regressions for each 2-digit sector:

$$y_{i,t} = \alpha + \eta_t + \delta_i + \beta_{M,s} m_{i,t} + \beta_{K,s} k_{i,t} + \varepsilon_{i,t}$$

There seems to be a high correlation among various estimates obtained using different assumptions on the production function. Table 5 displays these correlations. The translog specification simply includes second order terms for each of the inputs.[24]

Table 5: Pairwise Correlations Among TFP Estimates

	C-D, sect.	C-D, agg.	Translog, sect.	Translog, agg.
Cobb-Douglas, sectors	1			
Cobb-Douglas, agg.	0.904	1		
Translog, sectors	0.805	0.719	1	
Translog, agg.	0.899	0.979	0.736	1

[24]The translog version allowing different elasticities across 2-digit sectors is

$$y_{i,t} = \alpha + \eta_t + \delta_i + \beta_{M1,s} m_{i,t} + +\beta_{M2,s} m_{i,t}^2 + \beta_{K1,s} k_{i,t} + \beta_{K2,s} k_{i,t}^2 + \varepsilon_{i,t}$$

as before, TFP is simply the residual of this equation.

B Step Size Distributions for Results

Here I include a plot of the fixed distributions $F(n)$ discussed in section 3.3.2 under the two values of the radical innovation parameter ϕ used for the steady state comparisons in section 4.3. Recall that for a particular step size (size of quality improvement) $F(n) = c_0(n+m)^{-\phi}$. For the most backward firm at position $-\bar{m}$, the probability of getting a larger than one step improvement is about 10 percentage points higher in the baseline case than in the higher ϕ case.