Market Concentration and the Productivity Slowdown

Jane Olmstead-Rumsey*

June 15, 2022

Abstract

Since around 2000, U.S. aggregate productivity growth has slowed and product market concentration has risen. To explain these facts, I construct a measure of innovativeness based on patents that is comparable across firms and over time and show that small firms make innovations that are more incremental in the 2000s compared to the 1990s. I develop an endogenous growth model where the quality of new ideas is heterogeneous across firms to analyze the implications of this finding. I use a quantitative version of the model to infer changes to the structure of the U.S. economy between the 1990s and the 2000s. This analysis suggests that declining innovativeness of smaller firms can account for the bulk of the rise in market concentration and the productivity slowdown. Strategic changes in firms' R&D investment policies in response to the decreased likelihood of laggards making drastic improvements significantly amplify the productivity slowdown.

*London School of Economics and Federal Reserve Bank of Minneapolis, 90 Hennepin Ave., Minneapolis, MN. e-mail: jane.olmstead.rumsey@gmail.com. I especially thank Matthias Doepke, Guido Lorenzoni, and David Berger for their guidance and support. I thank the Minneapolis Fed, Heller-Hurwicz Institute, Chicago Fed, St. Louis Fed, Fed Board, and CREI for hosting me and providing invaluable feedback. I thank Ufuk Akcigit, Murat Alp Çelik, Nicolas Crouzet, Sergio Ocampo Díaz, Benjamin Jones, Kiminori Matsuyama, Martí Mestieri, Simon Mongey, and Dimitris Papanikolaou for providing comments at various stages. Finally, I am grateful to participants at numerous conferences and seminars for their comments and suggestions.

1 Introduction

After a boom in the 1990s, U.S. productivity growth declined in the early 2000s and industry leaders began capturing a larger share of sales. I illustrate a new mechanism that explains both these trends, along with patterns of increasing profitability, increasing productivity differences, and a declining rate of market leadership turnover in U.S. industries. This mechanism is the declining ability of laggard firms to catch and overtake market leaders through innovation.

To support the existence of this mechanism, I show that patent quality, one measure of innovativeness, fell sharply among smaller firms since 2000 after the 1990s boom. This fact is robust to using market-based measures of patent value or measures of the social value of patents using citations and is broad-based, spanning many sectors of the economy, though it is most pronounced in high tech sectors.

To understand how leaders (the largest firm in terms of sales in each industry) and laggards (other, smaller firms) respond to diminished opportunities for laggard firms to grow through innovation, and the effect on aggregate growth, I develop and estimate a general equilibrium, quality ladder model of directed innovation along the lines of Aghion et al. (2001). A continuum of industries are populated by two incumbent firms producing differentiated goods. Incumbent firms improve their varieties through innovation. Each industry also contains a competitive fringe of firms with the ability to imitate the incumbent that has the lower quality variety. Market concentration, measured as the market leader's share of industry sales, is high when the quality difference between the two incumbents' product varieties is large.

Unlike Aghion et al. (2001), the model accommodates the possibility that laggard firms have an "advantage of backwardness," allowing them to improve their variety more drastically than market leaders when they innovate. A model parameter governs the extent of laggard firms' advantage of backwardness. Other model parameters capture alternative reasons for rising market concentration and/or slowing productivity growth, that have been suggested such as slowing knowledge diffusion or entry rates, rising market power, and declining real interest rates.

To infer the relative importance of different changes to the U.S. economy in explaining these trends, I estimate the model parameters for two steady states to match data on concentration, productivity growth, the profit share, patent quality, the rate

of turnover in market leadership, aggregate R&D expenditures, and R&D expenditures at the firm level for the U.S. in two separate periods, the 1990s and the 2000s. This exercise suggests a dominant role for the parameter governing laggard firms' advantage of backwardness to explain rising concentration and slowing productivity growth compared to other explanations.

I use the model to explore the channels through which laggards' declining patent quality can explain trends in concentration and productivity growth. When laggards firms' advantage of backwardness declines, they respond to a lower chance of attaining market leadership by investing less in R&D. Facing a lower probability of being overtaken, leaders invest slightly more. Together, these decisions lead to larger average quality differences between leaders and laggards in steady state. Quality differences map directly to leaders' market shares and markups, so that sales concentration and the profit share of total output also rise. This change fully explains the observed rise in concentration between the 1990s and the 2000s and explains about 55 percent of the increase in the profit share. The dynamics of the transition from high to low laggard patent quality match the relative speeds of the productivity slowdown, which happened quickly, and the rise in concentration, which has been more gradual.

The source of endogenous growth in the model is quality improvements to the differentiated products that firms produce. Making laggards' quality improvements more incremental generates a productivity slowdown through two channels. One is direct: even if firms devoted the same resources to research and development, the economy would grow more slowly because average quality improvements are smaller. Second, there is a strategic effect that amplifies the productivity slowdown: because of the relocation of R&D activity towards market leaders, whose innovations tend to be more incremental, the economy grows even more slowly than before. A growth decomposition exercise finds that this latter force accounts for roughly three quarters of the productivity slowdown in the model. Quantitatively, the estimated decline in laggards firms' advantage of backwardness generates a productivity slowdown in the model of a similar magnitude to the slowdown observed in the U.S.[1]

[1]The fact that changing innovativeness alone can explain the entire productivity slowdown does not rule out other explanations, since there may be forces working to increase productivity growth that the model does not capture such as population growth, entry, improvements in human capital, and globalization.

Related Literature This paper contributes a novel mechanism to the large and growing literature linking trends in concentration, productivity growth, and business dynamism using models of endogenous growth. Several papers emphasize the increasing importance of intangible assets and information and communications technology (ICT) as a possible explanation (Aghion et al. (2019a), de Ridder (2021), Corhay, Kung, and Schmid (2020)), or a more general rise in fixed operating costs for large firms (Ghazi 2021). Non-technological explanations include demographic changes (Hopenhayn, Neira, and Singhania (2018), Jones (2020), Peters and Walsh (2021), Karahan, Pugsley, and Sahin (2019), Engbom (), Eggertsson, Mehrotra, and Robbins (2019), Bornstein (2021)) or declining real interest rates (Liu, Mian, and Sufi (2021), Chatterjee and Eyigungor (2020)). The most closely related explanation is the one in Akcigit and Ates (2020) and Akcigit and Ates (2019) that diffusion of knowledge from leaders to laggards is slowing down, either because of ICT and the increasing importance of data in firms' production processes or because of anti-competitive use of patents.

Rather than emphasizing particular features of information technology, I instead hypothesize that general purpose technologies (GPTs) may affect firm dynamics and market structure in addition to temporarily raising aggregate productivity growth. Past fluctuations in patent quality and productivity growth have been attributed to waves of innovation due to the arrival of new GPTs (Kelly et al. (2021); Kogan et al. (2017); Liu and Ma (2021)). Bresnahan and Trajtenberg (1995) define GPTs as new technologies are applicable in a wide range of sectors and exhibit innovational complementarities, meaning that they increase the productivity of downstream research and development efforts.[2] Given the new evidence presented here on heterogeneity in patent quality across firms and time, I argue that these innovational complementarities appear to be stronger for smaller firms than for market leaders.

Most neo-Schumpeterian growth models assume goods within sectors are perfect substitutes so that each sector has just one producer in each period (see Klette and Kortum (2004), Lentz and Mortensen (2008), Acemoglu and Cao (2015), and Akcigit and Kerr (2018) for leading examples).[3] Because of this, these models are not well-

[2]See Brynjolfsson, Rock, and Syverson (2021) for further discussion.

[3]A notable exception is Peters (2020), but he does not look at drivers of changes in concentration over time.

suited to address industry-level moments such as sales concentration. Introducing a duopoly (plus a competitive fringe) allows me to make unified predictions about market concentration at the industry level and firm-level innovation rates, and makes not only markups but also sales concentration within sectors an endogenous outcome of the innovation process.

The duopoly formulation also 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. Strands (ii) and (iii) typically rely on opposing assumptions. According to the literature on rising market power, incumbent firms exercise greater pricing power now than in the past and this is reflected in rising markups and profitability (de Loecker, Eeckhout, and Unger (2020), De Loecker, Eeckhout, and Mongey (2021), Barkai and Benzell (2018)). On the other hand, the literature on superstar firms 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. 2020). The model resolves this seeming contrast by demonstrating how markups can rise at the same time as there is reallocation to relatively more productive firms without any changes at all to consumer preferences or the aggregate production function. The model is also consistent with the finding of Kehrig and Vincent (2021) that being a superstar firm is a temporary rather than permanent status. In the model, the relative advantage of high value added firms grows in the 2000s and the average duration of these "shooting star" spells increases, but these firms are eventually displaced by competitors.

The model's industry structure with imperfect substitutes makes it possible to quantitatively compare explanations for increased markups and profits in recent years to the superstar firm hypothesis that greater price sensitivity has sparked reallocation to large, productive firms. Within the model, neither story matches the data as well as a decline in laggards' patent quality, though I show that the superstar firm experiment generates a productivity slowdown alongside rising concentration in the estimated model. To my knowledge, this is the first dynamic version of Autor et al. (2020) with endogenous productivity growth.

The finding that laggards' patent quality has declined since 2000 is consistent with Bloom et al. (2020), who show that despite increasing inputs (expenditures, workers) to R&D, outputs in terms of productivity improvements have declined. Anzoategui et al. (2019) identify a decline in R&D productivity using indirect inference in a DSGE model with endogenous productivity growth. Several papers have documented that laggard firms are less likely to overtake market leaders in recent years (Bessen et al. (2020), Pugsley, Sedlacek, and Sterk (2021), Andrews, Criscuolo, and Gal (2016)). This paper sheds more light on the channel through which this happens: I estimate a mild *decrease* in the cost per patent to explain rising expenditures on R&D over this period, but also a large decrease in the average contribution of a new patent to the value of the firm for laggard firms. Contemporaneous work by Cavenaile, Celik, and Tian (2020) estimates an endogenous growth model with incumbents and a competitive fringe with step by step innovations and finds that declining R&D productivity of small firms can explain a large share of the rise in concentration and the productivity slowdown. The advantage of allowing for patent quality heterogeneity and including new data on patent quality as a target for the estimation is that I can separately identify changing costs and changing output of R&D.

Finally, many papers have studied rising concentration and the productivity slowdown (Hall (2015), Syverson (2017)) in isolation from one another. Rising concentration is mainly a within-sector phenomenon (Hsieh and Rossi-Hansberg 2021) that is occurring at the national product market level rather than at the local level.[4] The finding that concentration is rising is robust to the inclusion of foreign firms (Covarrubias, Gutierrez, and Philippon 2019) or more sophisticated methods of identifying firms' competitors (Pellegrino (2021), using data from Hoberg and Phillips (2016)).[5]

A variety of explanations for rising sales concentration have been proposed, from the introduction of ICT that creates winner-take-all markets and enables the growth of superstar firms (Bessen (2020), Crouzet and Eberly (2018)), to excessive regulations that erect barriers to entry and create unnatural monopolies (Covarrubias, Gutierrez, and Philippon (2019)), to increased mergers and acquisitions activity, possibly due

[4] In fact Berger, Herkenhoff, and Mongey (2021) and Rossi-Hansberg, Sarte, and Trachter (2020) find evidence that local sales concentration has fallen over this period.

[5] Amiti and Heise (2021) do find that concentration has been stable in manufacturing industries once foreign firms are accounted for but do not expand their analysis to other industries.

to weak antitrust enforcement (Grullon, Larkin, and Michaely (2019)). This paper complements these hypotheses by contributing a novel mechanism that, according to the quantitative exercise, explains a large share of the increase.

2 Empirical Motivation

2.1 Market Concentration and Productivity Growth

Figure 1 plots the average market leader's share of total industry sales in Compustat and the total factor productivity growth rate.[6] Among U.S. public companies, market concentration has risen significantly since the late 1990s.[7] The average market leader's sales share within narrowly defined 4-digit Standard Industrial Classification (SIC) industries has risen from around 40% in the l990s to over 50% in 2017. Annual total factor productivity growth averaged about 1.6% between 1994 and 2003, but slowed to about 0.7% on average between 2004 and 2017.

According to the standard Olley and 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. Second, reallocation to the most productive firms (i.e. the covariance of sales share and productivity) could be slowing. Baqaee and Farhi (2020) show that within-firm growth has contributed very little to aggregate TFP growth since the late 1990s while allocative efficiency has risen, lending support to explanations focusing on the incentives of existing firms to increase productivity.

2.2 Trends in Patent Quality

Economists have long relied on patents as an observable proxy for innovativeness (Griliches (1998)). The most commonly used measure of patent quality, counting the number of forward citations a patent receives from future patents, shows substantial

[6] I focus the analysis on non-financial, non-agricultural firms. See Appendix A.1 for a detailed description of the Compustat sample. Total factor productivity data comes from Comin et al. (2021).

[7] See Grullon, Larkin, and Michaely (2019) and Council of Economic Advisers (2016) for overviews of trends in market concentration. More than 75% of U.S. industries have experienced an increase in the Herfindahl-Hirschman index.

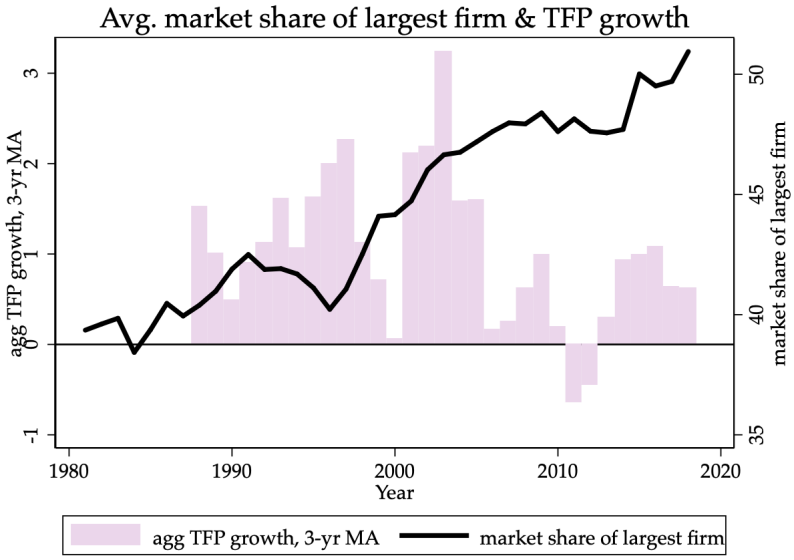

Figure 1: Average market share of largest firm (by sales) in 4-digit SIC industries from Compustat (weighted by industry sales); total factor productivity (TFP) estimates from Comin et al. (2021), three year moving average. See Appendix A.1 for a description of the Compustat sample.

heterogeneity in quality in the cross section of patents, with a few patents receiving many citations and most receiving none or just a few (Akcigit and Kerr 2018).

Recent evidence using alternative measures of patent quality also points to substantial changes in average quality over time. Kelly et al. (2018) create a text-based measure of patent quality, identifying "breakthrough" patents as those patents where the patent's text differs from the text of past patents but is similar to the text of future patents. This measure has the advantage of covering a longer time series (1860-present) than citation based measures (1940-present). Kelly et al. (2018) find that periods with high average patent quality coincide with the discovery of new general purpose technologies, including the ICT revolution in the 1990s, consistent with Bresnahan and Trajtenberg (1995)'s theory of "innovational complementarities" between general purpose technologies and inventions in other sectors of the economy.[8] The most recent wave of high patent quality driven by ICT began to subside in the late 1990s according to this measure (Figure A.1 in Appendix A.3).

[8]See Helpman (1998) and Aghion, Akcigit, and Howitt (2014) for reviews of the study of GPTs.

To explore heterogeneity in the decline in patent quality across firms, I use a measure of patent value from Kogan et al. (2017) that estimates the market value of all patents issued in the U.S. and assigned to public firms using firms' excess stock returns in a window around patent approval dates to infer the market value of the patent.[9] This measure has the advantage of capturing the private value of the patent to the firm, which determines firms' investment decisions in the model.[10]

In the model, firms make innovations that grow the quality of their product by a random amount. I use the dollar value estimates of Kogan et al. (2017) to construct a measure of each public firm's "patent stock" as the cumulative value of all past patents, intuitively corresponding to the current knowledge or quality embodied in the firm's product(s).[11,12] I define *patent quality* as the marginal contribution of a new patent to the total value of the firm's existing patent stock. From 1980 to 2017 this measure covers 1,305,813 patents issued to 5,693 different U.S. public firms. Figure 2 plots the average of this measure over time, splitting the sample into market leaders (largest firms by sales in 4-digit SIC industries) and followers (all other public firms).

Figure 2 illustrates two key facts for the subsequent analysis: (1) smaller firms have higher patent quality than market leaders on average; (2) smaller firms' patent quality rose from 1990 to 2000, but has declined significantly since 2000.

Fact (1) is related to the debate on the relative innovativeness of large versus small firms (see Akcigit and Kerr (2018)). Typically this debate centers on small startups versus large companies with more than 500 employees (more than 72% of observations in the sample of patenting firms in Compustat have more than 500 employees). I find that even among firms that are large relative to the entire firm size distribution, there are differences in patent quality by size (measured by sales)

[9] I use the data updated through 2020 from the paper's webpage: https://github.com/KPSS2017/Technological-Innovation-Resource-Allocation-and-Growth-Extended-Data.

[10] Kogan et al. (2017) show that this measure is strongly correlated with forward citation-weighted measures of patent quality at the patent level. Kelly et al. (2021) show that this measure is also correlated with the measure of patent quality based on the novelty of a patent's text.

[11] Construction details in Appendix A.2.

[12] Some depreciation can be applied to the patent stock. Applying rates in the 5-20% range used in Peters and Taylor (2017) increases the level of the estimated quality improvements but does not affect the magnitude of the slowdown or the differential decline between leaders and laggards.

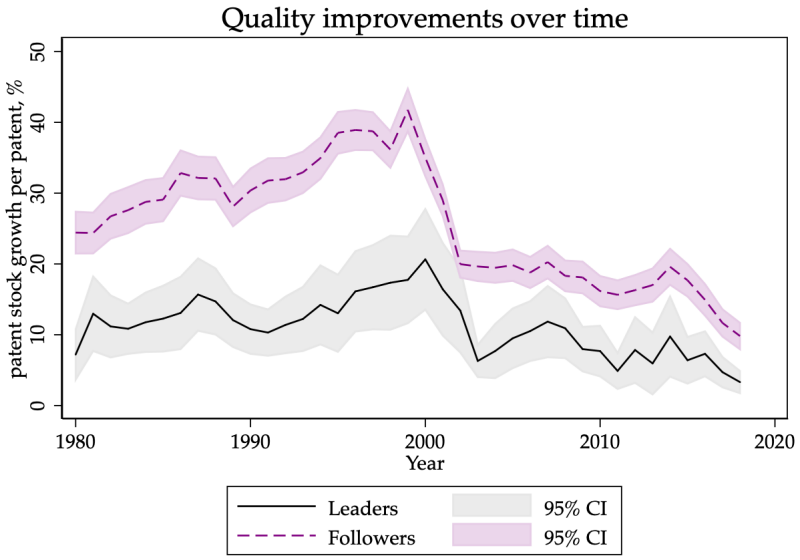

Figure 2: Contribution of average new patent to value of filing firm's existing stock of patents, using estimated patent values from Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries in a given year and followers are all other public firms in the sample. Three year moving averages. See Appendix A.1 for more details.

within industries. Managers at smaller firms tend to be more flexible and to be closer to both customers and researchers within the firm, enhancing their ability to allocate spending to more productive projects (Knott and Vieregger 2016). Rosen (1991) develops a model where large firms optimally focus on innovations that are complementary to their existing products to avoid Arrow (1962)'s replacement effect. Small firms instead fund the development of disruptive technologies, having more to gain in post-innovation rents from doing so.

One possible explanation for fact (2) is that general purpose technologies, or at least ICT, have more complementarity with the R&D investments of some types of firms than others (Jovanovic and Rousseau (2005) find that initial public offerings surge during GPT waves, for example). Smaller firms, with greater flexibility and more incentive to invest in riskier, disruptive ideas, may be better positioned to take advantage of the innovation opportunities associated with disruptive technologies. After the GPT has diffused through the economy, opportunities for disruption lessen

and laggards' improvements become more incremental.[13] Consistent with this idea, the pattern of boom and bust in patent quality is more pronounced in high tech sectors than in manufacturing, healthcare, or consumer goods, but the trend is present to some extent in all four categories (see Figure A.3).

The sharp decline in laggards' patent quality between 1999 to 2001 is worth exploring. The only significant change to U.S. patent law in the late 1990s was the American Inventor's Protection Act of 1999 to publish patent applications 18 months after they were filed. Previously, only approved patents were published. This change might deter inventors who thought their patent was unlikely to be approved from applying for fear that their idea would be published but they would not obtain the patent. In that case one would expect the patent approval rate to rise. In fact, the approval rate declined from about 70% in 1996 to 40% in 2005 (Carley, Hegde, and Marco 2015). Moreover, Graham and Hegde (2015) find that firms given the option to opt out of this pre-grant disclosure chose to do so less than 10% of the time.

The decline is likely not driven by the dot-com bubble. The same pattern appears in an analogous measure of patent quality based on citations (Figure A.2). It also seems not to be driven by the aging of public firms: this pattern appears even among firms that had been public at least 20 years when the patent was issued (Figure A.4). Nor is it driven by ideas being embodied in multiple patents in recent years: roughly the same pattern is present in the annual patent stock growth rather than the marginal contribution of each individual patent (Figure A.5).

2.3 Declining Dynamism and Leadership Turnover

In the model, innovations drive growth in market share at the expense of the firm's competitors. Figure 3 plots the fraction of U.S. industries with a new sales leader each year to measure the frequency with which smaller firms overtake the largest firm. This fraction has fallen from around 15% per year in the late 1990s to around 9% (see Bessen et al. (2020) for a detailed empirical analysis of this phenomenon in the U.S. and Andrews, Criscuolo, and Gal (2016) for a cross-country analysis).

Firm-level productivity data also shows that the "advantage of backwardness"

[13] Aum, Lee, and Shin (2018) find that the productivity boom from computerization had normalized by 2004, for example.

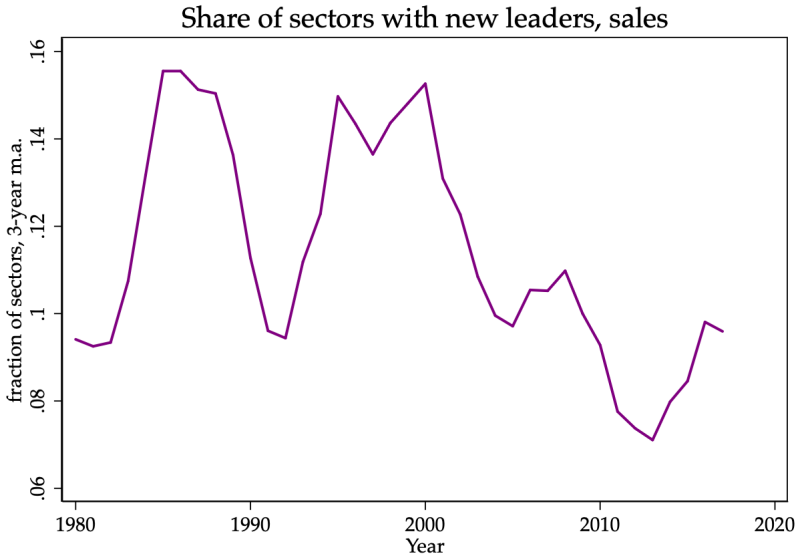

Figure 3: Share of 4-digit SIC sectors in Compustat with a new largest firm by sales in each year, three year moving average. See Appendix A.1 describing the Compustat sample.

has fallen relative to the 1990s. Andrews, Criscuolo, and Gal (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 most productive competitor has decreased over the 2000s, suggesting that distance to the productivity frontier is becoming a less important predictor of future productivity growth. Decker et al. (2016) also find that the right skewness of the firm-level productivity growth distribution in the U.S. has declined over this period.

3 Model

To capture the effect of declining innovativeness of laggards firms, I develop a model along the lines of Aghion et al. (2001) but building on models with heterogeneous patent quality rather than step by step innovations (Akcigit, Ates, and Impullitti (2018), Acemoglu et al. (2018), Akcigit and Kerr (2018)). Relative to Aghion et al. (2001), I also introduce a competitive fringe of firms in each sector that constrains the pricing behavior of the incumbents in order to match levels of concentration in the data. Markups are endogenous and each sector's level of sales concentration varies

over time as the result of innovation. The model is of a closed economy in continuous time. There are three types agents: a representative household, a representative final good firm, and firms producing intermediate goods.

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:

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

where ρ is the discount rate, ψ is the inverse intertemporal elasticity of substitution, C_t is consumption at time t, W_t is the nominal wage rate, and P_t is the price of the consumption good C_t. Households' labor supply L will be normalized to 1 and there is no population growth. Households own all firms, and the total assets in the economy A_t are:

$$A_t = \int_0^1 \sum_{i=1}^{2} \left(V_{ijt} + V_{ijt}^e\right) dj,$$

where V_{ijt} is the value of an incumbent intermediate good firm i in sector j at time t and V_{ijt}^e is the value of an entrant that can displace firm i in sector j at time t. These value functions are explained in greater detail in section 3.3. r_t is the rate of return on the portfolio of firms. On a balanced growth path with constant growth rate of output g this yields the standard Euler equation $r = g\psi + \rho$.

3.2 Final Good 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 operates a constant return to scale technology:

$$Y_t = \frac{1}{1-\beta}\left(\int_0^1 K_{jt}^{1-\beta} dj\right) L^\beta, \tag{1}$$

where K_{jt} is a composite of two products produced by sector j described below. β determines the elasticity of substitution across sectors ($\frac{1}{\beta}$) and the labor share.[14] The final good firm's problem of hiring sector composite goods K_{jt} for $j \in [0,1]$ and labor is:

$$\max_{K_{jt},L} P_t \frac{1}{1-\beta} \left(\int_0^1 K_{jt}^{1-\beta} dj \right) L^{\beta} - \int_0^1 P_{jt} K_{jt} dj - W_t L.$$

The first order condition for sector j's composite good given sector j's composite price index P_{jt} yields the demand for sector j's good:

$$K_{jt} = \left(\frac{P_{jt}}{P_t} \right)^{-\frac{1}{\beta}} L,$$

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

$$\beta \frac{Y_t}{L} = \frac{W_t}{P_t}.$$

To derive the demand curve for each intermediate good producer i in sector j we need to define the sector composite goods K_{jt} explicitly:

$$K_{jt} = \left((q_{1jt} k_{1jt})^{\frac{\epsilon-1}{\epsilon}} + (q_{2jt} k_{2jt})^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}, \tag{2}$$

where q_{ijt} is the quality of firm i's product at time t (equivalently firm i's productivity) and k_{ijt} is the output of firm i purchased by the final good producer.[15] The elasticity of substitution between product varieties in the same sector is ϵ.

The first order condition for the final goods firm's problem yields the following demand curve for firm i in sector j's output:

$$k_{ijt} = q_{ijt}^{\epsilon-1} \left(\frac{p_{ijt}}{P_{jt}} \right)^{-\epsilon} \left(\frac{P_{jt}}{P_t} \right)^{-\frac{1}{\beta}} L. \tag{3}$$

That is, demand is increasing in the firm's quality, decreasing in its price relative to the sector j price index, and decreasing in the sector's price index relative to the price index in the economy as a whole.

[14] In Section 4.1 I discuss how a standard calibration of the labor share implies a cross-sector elasticity of substitution that is consistent with estimates of this parameter in the data.

[15] I use quality and productivity interchangeably because final output is homogeneous of degree one in either the qualities or the quantities of the intermediate goods firms' products and labor.

3.3 Intermediate Goods Producers

Each intermediate good sector features competition between two large incumbent firms with differentiated products and access to an R&D technology, plus a competitive fringe that constrains the price-setting of the incumbents. Incumbents are occasionally hit with exit shocks that cause them to be replaced by a new firm.

3.3.1 Production and Price Setting

Production Intermediate goods producers purchase final goods to transform them into differentiated intermediate goods. Each unit of intermediate output requires $\eta < 1$ units of the final good to produce. There are no other inputs to intermediate good production.

Competitive fringe Each industry has a competitive fringe of firms that can produce a perfect substitute to the lower quality variety at marginal cost η. I call the incumbent firm with lower quality the *follower*, or laggard, and the incumbent firm with higher quality the *leader*. When $q_{1jt} = q_{2jt}$, the fringe can produce perfect substitutes to both varieties. One way to micro-found this assumption is to introduce a cost to filing and maintaining 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 in duopoly competition without the fringe, would be willing to pay for exclusive rights to its variety. Intuitively, this means that sectors in the model feature a high quality variety like a brand name product and competition among other firms to produce a generic version. The competitive fringe firms do not have access to an innovation technology.

Assuming the presence of a competitive fringe is not necessary to solve the model, but makes it possible to match the average level of sales concentration across sectors in the data and generates plausible predictions for profit shares as a function of market shares (see Appendix A.5.) I solve a version of the model without the competitive fringe in Appendix B.5 and replicate the main exercise in this setting. The main results in Section 4 are qualitatively unchanged.

Price setting Firms set prices a la Bertrand at each instant t. The presence of the fringe implies the follower must set its price $p_{ijt} = \eta$.[16] The leader chooses its

[16] I resolve the indeterminacy of which firm(s) produces the lower quality variety in equilibrium

price as a best response to the follower's price.

Omitting the subscript t, the static problem of leader i in sector j is:

$$\max_{p_{ij}} p_{ij}k_{ij} - \eta k_{ij},$$

subject to the demand:

$$k_{ij} = q_{ij}^{\epsilon-1}\left(\frac{p_{ij}}{P_j}\right)^{-\epsilon}\left(\frac{P_j}{P}\right)^{-\frac{1}{\beta}} L,$$

where

$$P_j = \left(\sum_{i=1}^{2} q_{ij}^{\epsilon-1} p_{ij}^{1-\epsilon}\right)^{\frac{1}{1-\epsilon}}$$

is sector j's price index.

Let $s_{ij} = \frac{p_{ij}k_{ij}}{\sum_{i=1}^{2} p_{ij}k_{ij}}$ denote firm i's market share in sector j. Then the optimal pricing policy for the market leader is:

$$p_{ij} = \frac{\epsilon - (\epsilon - \frac{1}{\beta})s_{ij}}{\epsilon - (\epsilon - \frac{1}{\beta})s_{ij} - 1}\eta. \tag{4}$$

The optimal price is the standard two-layered constant elasticity of demand (nested CES) solution: a variable markup rising in market share.

3.3.2 Innovation

Incumbent intermediate goods producers have access to a research and development technology that allows them to choose an amount of research spending R_{ijt} of the final good to maximize the discounted sum of expected future profits. The decision to model R&D as a process of own-product quality improvement by incumbents is consistent with the evidence in Garcia-Macia, Hsieh, and Klenow (2019) that: (i) incumbents are responsible for most employment growth in the U.S., and this share has increased in recent years; (ii) growth mainly occurs through quality improvements rather than new varieties; (iii) creative destruction by entrants and incumbents over other firms' varieties accounted for less than 25% of employment growth from 2003-2013, consistent with earlier evidence from Bartelsman and Doms (2000).

by having the incumbent capture all sales of the lower quality variety so the fringe is not active in equilibrium.

Innovations arrive randomly at Poisson rate x_{ijt} which depends on research spending according to the function:

$$x_{ijt} = \left(\frac{\gamma R_{ijt}}{\alpha}\right)^{\frac{1}{\gamma}} q_{ijt}^{\frac{1-\frac{1}{\beta}}{\gamma}}.$$

Since $\beta < 1$, with higher quality more R&D is needed to achieve the same arrival rate x. γ and α are R&D technology parameters.

Innovations improve the quality of the incumbent firm's variety.[17] Conditional on innovating the size of the quality improvement is random. Formally, conditional on innovating,

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

where $\lambda > 1$ is some minimum quality improvement and $n_{ijt} \in \mathbb{N}$ is a random variable. Note that each competitor improves over their own quality when they innovate, rather than over the quality frontier.[18] Initial qualities of all firms at $t = 0$ are normalized to 1. Let $N_{ijt} = \int_0^t n_{ijs} ds$ denote the total number of λ step improvements over a product line i since the beginning of time. The *technology gap* m_{ijt} from firm 1 in sector j's perspective at time t is defined as:

$$\frac{q_{1jt}}{q_{2jt}} = \frac{\lambda^{N_{1jt}}}{\lambda^{N_{2jt}}} \equiv \lambda^{m_{1jt}}.$$

Given λ, m_{ijt} parameterizes the relative qualities of the two firms within sector j from firm $i \in \{1, 2\}$'s perspective, representing the number of λ steps ahead or behind its competitor firm i is. m_{ijt} turns out to be the only payoff relevant state variable for the incumbent firms. For tractability I impose a maximal technology gap $\bar{m}$, but in calibrating the model I will set the parameters so that this maximal gap rarely occurs in steady state. I assume that the only knowledge spillover between incumbents in the model occurs when a firm at the maximal gap innovates. In that case, both the innovating firm and its competitor's quality increase by the factor λ, keeping the technology gap unchanged but raising the absolute quality of the sector composite good.

[17] 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.

[18] Luttmer (2007) provides an additional rationale for this assumption: entrants are usually small and enter far from the productivity frontier, implying that imitation of others is difficult.

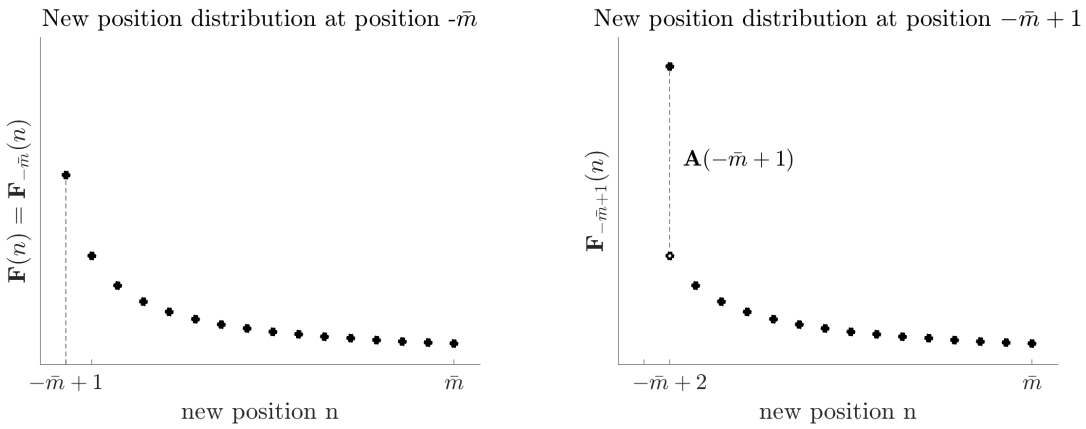

Figure 4: Examples of new position distributions for positions $-\bar{m}$ and $-\bar{m}+1$.

The distribution of possible quality improvements depends on the firm's current technology gap, consistent with the evidence in section 2 that patent quality varies between market leaders and laggards. Because of the subsequent result that firm strategies depend on technology gaps, it is useful to formulate these distributions as though firms draw a new position in technology gap space $n \in \{-\bar{m}, \ldots, \bar{m}\}$ when they innovate, rather than an absolute number of λ steps. Given n and the technology gap m the number of steps can then be derived as $n - m$.

As in Akcigit, Ates, and Impullitti (2018), I assume there exists a fixed distribution $\mathbb{F}(n) \equiv c_0(n + \bar{m})^{-\phi}$ for all $n \in \{-\bar{m}+1, \ldots, \bar{m}\}$ that applies to firms that are the furthest possible distance behind their competitor and describes the probability that they move to each position in technology gap space. An example is shown in the left panel of Figure 4. The shape parameter ϕ is critical in the model and determines the speed of catchup by increasing or decreasing the relative probability of larger innovations. A higher ϕ means a lower probability of these "radical" improvements.[19,20] c_0 is simply a shifter to ensure $\sum_n \mathbb{F}(n) = 1$.

Given this fixed distribution for the most laggard firm, the new position distribu-

[19] As noted by Akcigit, Ates, and Impullitti (2018), this formulation converges to the less general step-by-step model as $\phi \to \infty$.

[20] The use of "radical innovation" in this paper to describe a relatively large quality improvement differs from some other papers in the literature such as Acemoglu and Cao (2015) who use "radical innovation" to refer to an entrant replacing an incumbent.

tion for 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 4 for a firm at gap $-\bar{m}+1$. All the mass of the fixed distribution on positions lower than the current position m is put on one-step ahead improvements. This formulation captures the feature that laggard firms make larger improvements than leaders on average.

3.3.3 Entry and Exit

Incumbents face a constant exit risk δ_e. If an incumbent is hit with this shock they are replaced by an entrant that takes over the product line with the same quality level (and thus technology gap to the other incumbent in the sector) as the incumbent it replaces. This shock captures reasons why incumbents may exit or be displaced by entrants that are not related to the incumbent firms' innovations such as adverse financial shocks, negative taste shocks for the incumbent's brand, expiration of the incumbent's patent or knowledge diffusion as in Akcigit and Ates (2020), or cost shocks to specific inputs used by the incumbent.

3.3.4 Intermediate Goods Firms Value Functions

Turning to the firm value functions, I will show that the technology gap $m \in \{-\bar{m}, \ldots, \bar{m}\}$ is sufficient to describe the firms' pricing and innovation strategies, and that firm values scale in a particular function of their current product quality q_{ijt}.

The proof that pricing decisions and market shares depend only on m_{ijt} (and not on the level of quality q_{ijt}) is in Appendix B.1. The flow profits of an incumbent, denoting the optimal price of the leader at technology gap m as $p(m)$, are:

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

Using equation 3 for k_{ijt} and the definition of the sector price index:

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

For the dynamic problem, I use a guess and verify method to verify that firms' strategies depend only on m and that firm values scale in some function of q_{ijt}. Given an interest rate r_t, the value function of a firm with gap m to its competitor and quality level q_{ijt} can be written:

$$\begin{aligned} r_t V_{mt}(q_{ijt}) - \dot{V}_{mt}(q_{ijt}) = \max_{x_{mt}} \{ & \pi(m, q_{ijt}) - \alpha \frac{(x_{mt})^\gamma}{\gamma} q_{ijt}^{\frac{1}{\beta}-1} \\ & + x_{mt} \sum_{n_t=m+1}^{\bar{m}} \mathbb{F}_m(n_t)[V_{nt}(\lambda^{n_t-m} q_{ijt}) - V_{mt}(q_{ijt})] \\ & + x_{(-m)t} \sum_{n_t=-m+1}^{\bar{m}} \mathbb{F}_{-m}(n_t)[V_{(-n)t}(q_{ijt}) - V_{mt}(q_{ijt})] \\ & + \delta_e(0 - V_{mt}(q_{ijt})\}. \end{aligned} \tag{5}$$

The firm chooses the arrival rate of innovations x_{mt}. The first line shows the flow profits and the R&D cost R_{ijt} given the choice of x_{mt}. The second line gives the probability that the firm innovates and sums over the possible states the firm could move to using the new position distribution and the firm's new value function with a higher level of product quality *and* a higher *relative* quality compared to its industry rival. The third line denotes the chance that the firm's rival innovates and the change in the firm's value because its relative quality falls. The final line shows the chance an entrant displaces the incumbent. The slightly altered equations for firms at the minimum and maximum gaps due to knowledge spillovers are given in Appendix B.2.

A guess and verify approach verifies that $V_{mt}(q_{ijt}) = v_{mt} q_{ijt}^{\frac{1}{\beta}-1}$. Thus one can focus on a Markov perfect equilibrium where firms' strategies depend only on the payoff-relevant state variable m which characterizes the technology gap between incumbents.

The firm's optimal innovation rate x_{mt} is the solution to the first order condition

of equation (5), which gives:

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

Intuitively, firms choose a higher arrival rate of innovations when the cost of R&D α is low, and when the expected gain from innovating is high, captured by the probability of moving to different positions in technology gap space upon innovating $\mathbb{F}_m(n)$, the value v_n of being at gap n, and the minimum size of quality improvements λ. All else equal, greater expected innovativeness of laggards (higher probability that they catch up to or overtake the leader), encourages more innovation by laggard firms. However, the v_n terms also capture the probability of being displaced in the future, so these values are endogenously determined along with the chance of displacement by rivals due to innovation or the chance of being hit with an exit shock δ_e. At t, the value of a potential entrant in product line i in sector j is simply $V^e_{ijt} = \delta_e V_{ijt}$.

3.4 Equilibrium Output

Plugging in the intermediate goods firms' pricing decisions yields the following expression for final output Y_t, derived in Appendix B.3:

$$Y_t = \frac{1}{2} \frac{L}{1-\beta} P^{\frac{1-\beta}{\beta}} \sum_{m=-\bar{m}}^{\bar{m}} Q_{mt}, \tag{6}$$

where Q_{mt} is defined as:

$$\begin{aligned} Q_{m,t} &= \int_0^1 \left(q_{it}^{\epsilon-1} p(m)^{1-\epsilon} + q_{-it}^{\epsilon-1} p(-m)^{1-\epsilon} \right)^{-\frac{(1-\beta)}{\beta(1-\epsilon)}} \mathbb{1}_{\{i \in \mu_{mt}\}} di \\ &= (p(m)^{1-\epsilon} + (\lambda^{-m})^{\epsilon-1} p(-m)^{1-\epsilon})^{\frac{1-\beta}{\beta(\epsilon-1)}} \int_0^1 q_{i,t}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i \in \mu_{mt}\}} di. \end{aligned} \tag{7}$$

Here μ_{mt} is the measure of firms at each technology gap m at time t (normalizing the measure of firms to one) and Q_{mt} is a particular index of the qualities of all firms at gap m. The change in output between t and $t + dt$ will therefore depend on the changes $\dot{Q}_{mt}$ 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

$\mathbb{F}(n)$. The term $(p(m)^{1-\epsilon} + (\lambda^{-m})^{\epsilon-1}p(-m)^{1-\epsilon})^{\frac{1-\beta}{\beta(\epsilon-1)}}$ weights the change in qualities of firms at gap m depending on the prices set by firms at gap m and $-m$, capturing static distortions from firms' markups. Note that entry and exit are not a source of growth in the model because they have no impact on the qualities of the intermediate goods in the economy or on markups. The final component determining output will be the measure of firms at each technology gap μ_{mt} that is itself an endogenous object. The next section describes how to solve for the measures μ_{mt}.

3.5 Distribution Over Technology Gaps

Firms move to gap n through innovation from a lower technology gap m, or because their competitor innovates to gap $-n$. The distributions $\mathbb{F}_m(n)$ and $\mathbb{F}_{-m}(-n)$ respectively determine these probabilities, combined with the innovation efforts of firms at m and $-m$, for all $m < n$ and $-m < -n$. The outflows from gap n are due to a firm at n or its competitor at $-n$ innovating. Putting this together into the Kolmogorov forward equations for the evolution of the mass of firms at each gap:

$$\dot{\mu}_{nt} = \sum_{m=-\bar{m}}^{n-1} x_m \mathbb{F}_m(n)\mu_{mt} + \sum_{m=n+1}^{\bar{m}} x_{-m}\mathbb{F}_{-m}(-n)\mu_{mt} - (x_n + x_{-n})\mu_{nt}. \tag{8}$$

The highest and lowest gaps are special cases because of spillovers: if the firm at the highest gap innovates both firms remain at the same gap:

$$\dot{\mu}_{-\bar{m}t} = \sum_{m=-\bar{m}+1}^{\bar{m}} x_{-m}\mathbb{F}_{-m}(\bar{m})\mu_{mt} - x_{-\bar{m}}\mu_{-\bar{m}t} \tag{9}$$

$$\dot{\mu}_{\bar{m}t} = \sum_{m=-\bar{m}}^{\bar{m}-1} x_m \mathbb{F}_m(\bar{m})\mu_{mt} - x_{-\bar{m}}\mu_{\bar{m}t}. \tag{10}$$

On a balanced growth path, $\mu_{mt} = \mu_m$ for all m, t. Replacing the left sides with zero change in equilibrium and the measures on the right sides with the constants μ_n, μ_m defines a system of $2\bar{m} + 1$ equations in $2\bar{m} + 1$ unknowns that determine the steady state distribution of firms over possible technology gaps. There are two additional restrictions on the solution to this system: (1) for each firm at m there is a firm at $-m$ (that is, the stationary distribution is symmetric); (2) I impose the restriction that the measure of all incumbent firms sums to one.

3.6 Output Growth

Differentiating equation 6 with respect to time yields the following expression for the growth rate:

$$\frac{\dot{Y}_t}{Y_t} = g_{Yt} = \frac{1}{2}\frac{1}{1-\beta}\sum_{m=-\bar{m}}^{\bar{m}}\frac{\dot{Q}_{mt}}{Y_t}.$$

It's useful to define:

$$\tilde{Q}_{mt} = \int_0^1 q_{m,t,i}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i\in\mu_{mt}\}} di \tag{11}$$

so that:

$$g_{Yt} = \frac{1}{2}\frac{1}{1-\beta}\sum_{m=-\bar{m}}^{\bar{m}}\left(p(m)^{1-\epsilon} + (\lambda^{-m})^{\epsilon-1}p(-m)^{1-\epsilon}\right)^{\frac{1-\beta}{\beta(\epsilon-1)}}\frac{\dot{\tilde{Q}}_{mt}}{Y_t}.$$

The subsequent analysis focuses on a balanced growth path where $\frac{\dot{\tilde{Q}}_{mt}}{Y_t}$ is constant for all m. On this balanced growth path consumption and output grow at a constant growth rate g and the mass of firms at each technology gap μ_m is constant. In general it is not possible to solve for this growth rate in closed form, but for a given set of model parameters it is possible to check the existence and uniqueness of such a balanced growth path and find the value of g as the solution to a system of equations. A derivation of these results is provided in Appendix B.4.

3.7 Equilibrium Definition

Let $R_t = \int_0^1 \sum_{i=1}^2 R_{ijt}dj$ denote total research and development spending by incumbents, 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. A Markov perfect equilibrium is an allocation $\{k_{ijt}, K_t, x_{ijt}, R_t, Y_t, C_t, L, \mu_{mt}, Q_{mt}, A_t\}_{i\in\{1,2\},j\in[0,1],m\in[-\bar{m},\bar{m}]}^{t\in(0,\infty)}$ 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. Households choose C_t and A_t to solve the problem in section 3.1.

2. Final goods firms solve their problem to hire labor L and buy intermediate goods k_{ijt} optimally given the problem in section 3.2.

3. Intermediate good firms choose p_{ijt} and x_{ijt} to solve their innovation and price-setting problems described in section 3.3.

4. The final goods market clears: $Y_t = C_t + R_t + K_t$.

5. The asset market clears, pinning down r_t via the Euler equation.

6. Labor market clears, pinning down W_t from the final good firm's problem.

7. μ_{mt} and Q_{mt} are consistent with firms' choices of x_{ijt}.

4 Model Estimation

I estimate an initial steady state for the model by matching various moments for the U.S. economy in the period of high patent quality between 1994 and 2003 ("1990s") using data on U.S. public firms from Compustat as well as aggregate moments. Using this initial calibration I describe firms' pricing and innovation strategies to develop intuition about the model. I then re-estimate the model parameters for 2004-2017 ("2000s") to infer changes to the economy between these two periods.

4.1 Baseline Calibration for the 1990s

Four parameters are calibrated outside the model using standard values. The inverse intertemporal elasticity of substitution ψ is set to 1. The labor share, β, is set to 0.6. This implies an elasticity of substitution across sectors of $\frac{1}{\beta} = \frac{5}{3}$, within the range of upper-level elasticities of substitution estimated in Hobjin and Nechio (2019). The curvature of the R&D cost function, γ, is calibrated outside to match the empirical evidence on the elasticity of patenting to R&D expenditures, discussed in Acemoglu et al. (2018). The maximal technology gap $\bar{m}$ is set to 16.

The rest of the parameters for the baseline model, shown in Table 1, are estimated using a simulated method of moments approach described in Appendix C.2 to match targets for the 1990s equilibrium. The targets are listed in Table 2. The data sources and computation methods for the data moments are given in Appendix A.2. Appendix C.1 describes the solution method for finding the model steady state.

	Value	Meaning/source
ψ	1	Inverse intertemporal elasticity of substitution
ρ	**0.0147**	Rate of time preference (annual)
β	0.6	Labor share/Nechio & Hobijn (2019)
ϵ	**4.203**	Elasticity of substitution within sectors
η	**0.628**	Marginal cost of intermediate producers
δ_e	**0.096**	Exogenous entry/exit rate (annual)
λ	**1.053**	Min. qual. improvement
γ	2	Curvature of R&D function
α	**4.826**	R&D cost parameter
$\bar{m}$	16	Maximum number of steps ahead
ϕ	**0.789**	Shape of patent quality distribution

Table 1: Model parameters, 1990s steady state. Estimated parameters in bold. Estimation details in the text and in Appendix C.2.

The targets include the main phenomena of interest: aggregate productivity growth, average market leader's share of industry sales, the profit share of total output, average patent quality, and the rate of leadership turnover from either entry or being overtaken by an incumbent rival. In addition to average patent quality across all firms, I include the average patent quality of market leaders and followers to help identify λ and ϕ separately. The other two moments, R&D as a share of output and R&D as a share of sales at the firm level, are included to help discipline the R&D cost parameter and the discount rate.

The model performs well in fitting the data, particularly for productivity growth and concentration. Intuitively, the minimum step size λ and ϕ govern the average patent quality, with λ acting as a vertical shift in patent quality for all types of firms and ϕ shifting the probability that laggards make drastic or incremental improvements, holding patent quality of leaders fixed. The R&D cost parameter α influences the amount all firms spend on R&D and helps match aggregate expenditures as a share of output and R&D as a share of firms' sales. The entry/exit shock δ_e helps match leadership turnover. One issue with the model fit is for R&D as a share of sales at the firm level. This can be attributed to the fact that productivity growth

	Data	Model
Avg. TFP growth, %	1.57	1.57
Avg. leader market share, %	43.23	43.23
R&D share of GDP, %	1.80	1.80
Profit share of GDP, %	5.24	5.69
Avg. R&D/sales, %	3.75	4.63
Avg. patent stock growth per patent, %	22.90	21.68
Avg. patent stock growth per patent, followers, %	24.86	33.72
Avg. patent stock growth per patent, leaders, %	10.00	9.99
Avg. leadership turnover, %	13.74	13.93

Table 2: Model fit for targeted moments from estimation of 7 parameters for the 1990s. Estimation details in the text and in Appendix C.2. The method for computing each moment in the data is described in Appendix A.2.

is purely due to R&D in the model, whereas in the reality productivity may improve for other reasons, like management practices or improved human capital.

The estimated parameters are reasonable: a discount rate ρ of 1.5% annually implies a real interest rate in the model of 3%. An elasticity of substitution ϵ of 4.2 results in an (unweighted) average markup of 1.23, in line with the evidence summarized in Mongey (2021), particularly de Loecker, Eeckhout, and Unger (2020), suggesting that average markups for U.S. public firms in the 1990s ranged from 1.2 to 1.3. The entry/exit rate of about 9.6% per year is in line with entry and exit rates for the U.S. reported by Decker et al. (2016). The model also matches non-targeted heterogeneity in R&D intensity (R&D as a share of sales). The ratio of followers' to leaders' R&D intensity in the data is 2.5 and in the model it is 3.2. Section 4.4 describes the model's fit for additional non-targeted moments.

4.2 Properties of the Baseline Model

The markups and market shares for quality leaders as a function of the leader's technology gap under the parameterization of the model in Table 1 are plotted in Figure 5. The leader's price $p(m)$ rises as the technology gap widens (that is, as the

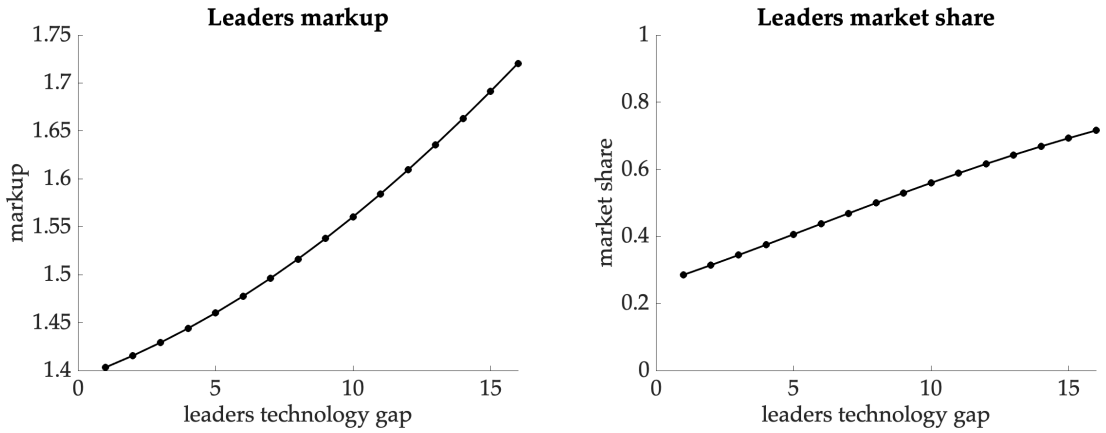

Figure 5: Leader's markup (price over marginal cost) and market share as a function of the leader's technology gap m in the model using the 1990s calibration in Table 1.

leader's relatively quality improves). The leader's market share rises from around 28% when the leader is one step ahead (when the leader's quality is 5.3% higher than the laggard's) to 72% of the market at the maximum 16 steps ahead (when the leader's quality is 130% higher than the laggard's). The follower, which sets price equal to marginal cost because of the presence of the fringe, has a large market share due to its relatively low price, and its market share is increasing in its relative quality.

The competitive fringe assumption also plays a role in determining the shape of the innovation policy as a function of technology gaps shown in Figure 6a, specifically the hump shape. This shape has been suggested theoretically in the work of Harris and Vickers (1987), Aghion et al. (2001), and Akcigit, Ates, and Impullitti (2018), and found in a variety of studies including Aghion et al. (2005), Aghion et al. (2018), Aghion et al. (2019b), and Zhang (2018). The hump shape appears in this model because the competitive fringe assumption means that the greatest incremental gain in flow profits comes from obtaining quality leadership (and thus escaping competition with the fringe), so the arrival rate of innovations is highest when the two firms have equal quality. In Appendix B.5 I show that the mechanism and main results are qualitatively unchanged when competitive fringe is eliminated.

Finally, Figure 6b shows the stationary distribution of sectors over the market leader's technology gap. Because followers in this equilibrium innovate more frequently than leaders and have a high chance of catching their competitor when they

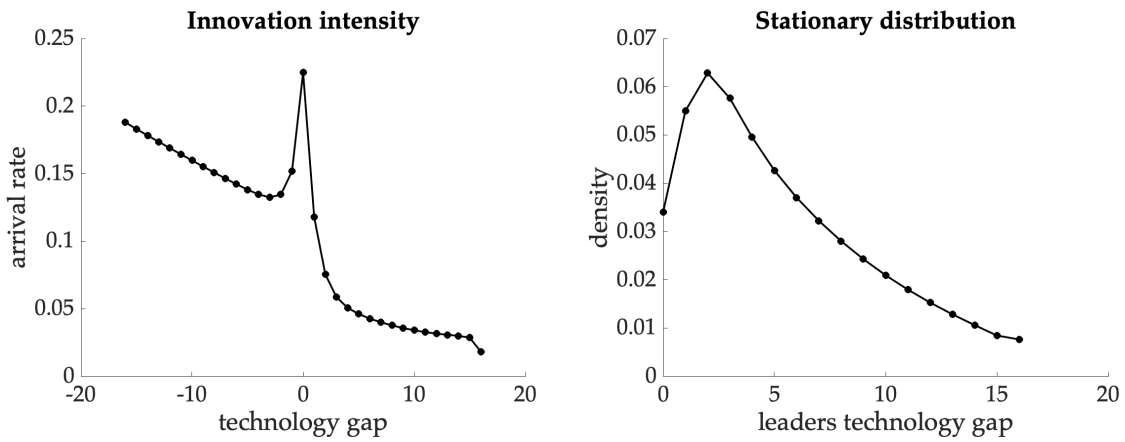

(a) Innovation policies **(b)** Stationary distribution

Figure 6: Innovation policies x as a function of the technology gap for all firms (panel a) and the stationary distribution of sectors over the technology gap of the leader (panel b) using the 1990s calibration in Table 1.

do, there is a high rate of turnover in market leadership and technology gaps do not grow large on average. Most sectors feature a leader only a few steps ahead of its rival, but there is a right tail of sectors with a dominant "superstar" that has much higher quality and thus captures a large share of industry sales.

4.3 Re-estimation for the 2000s

Re-estimating the model for the 2000s uses the model to infer the role of different channels suggested in the literature to explain changes in concentration and productivity growth and compare the strength of these other channels to the strength of declining laggard patent quality.

Changing the discount rate ρ captures the interest rate channel proposed by Liu, Mian, and Sufi (2021).[21] A decrease in entry and exit shocks δ_e can capture declining knowledge diffusion from incumbent firms to new entrants or rising entry costs (Akcigit and Ates (2020), Corhay, Kung, and Schmid (2020)). An increase in the research cost parameter α implies that more R&D spending is needed to achieve the

[21] For the estimated model, even at very low interest rates (for example 0.1%) the strategic effects described in Liu, Mian, and Sufi (2021) do not dominate the first-order effect of lowering the interest rate on firms' R&D. See Goldberg, Lopez-Salido, and Chikis (2021) for further discussion.

	Data	Model
Avg. TFP growth, %	0.66	0.67
Avg. leader market share, %	48.12	48.14
R&D share of GDP, %	1.89	1.58
Profit share of GDP, %	6.61	7.49
Avg. R&D/sales, %	4.06	4.14
Avg. patent stock growth per patent, %	11.22	10.75
Avg. patent stock growth per patent, followers, %	12.23	14.58
Avg. patent stock growth per patent, leaders, %	5.17	6.94
Avg. leadership turnover, %	9.28	9.44

Table 3: Model fit for targeted moments from estimation of 7 parameters for the 2000s. Estimation details in the text and in Appendix C.2. The method for computing each moment in the data is described in Appendix A.2.

same arrival rate of innovations, capturing the cost side of the hypothesis of Bloom et al. (2020) that ideas are getting harder to find. A decrease in the elasticity of substitution within sectors ϵ captures increased market power over the leader's variety, in line with Jones and Philippon (2016). On the other hand, an increase in ϵ captures the superstar firm hypothesis of Autor et al. (2020) that competitive pressures within industries have risen, causing the most productive firms to capture a larger share of total industry sales. Finally, changes in ϕ govern the expected patent quality for different types of firms by changing the distributions $\mathbb{F}_m(n)$. Changing ϕ represents the research *output* side of Bloom et al. (2020)'s hypothesis, capturing the possibility that the quality of new ideas, particularly for laggard firms, is falling.

Table 3 shows the targeted moments and model fit for the 2000s estimation. In the data, productivity growth slowed substantially compared to 1994-2003, while the average market leader's sales share grew by about five percentage points. Both aggregate and firm level R&D expenditures grew, as noted by Bloom et al. (2020). As discussed in detail in section 2, patent quality and leadership turnover declined. The estimation has some trouble matching the decline in the growth rate alongside an increase in R&D expenditure, but otherwise performs well.

Table 4 compares the estimated parameters for the two steady states. The house-

holds' discount rate declines in the 2000s. Consistent with Decker et al. (2016), the entry rate of new firms declines (alternately, incumbents are less likely to be displaced, consistent with the hypothesis of Akcigit and Ates (2019) that the rate of knowledge diffusion is slowing down). To match the fact that R&D expenditures as a share of GDP rose between the 1990s and the 2000s, the cost α of performing R&D declines. However, the expected *output* of R&D (patent stock growth per patent) conditional on innovating declines substantially due to the decrease in the probability of radical innovations, driven by the substantial increase in ϕ. The marginal cost of the intermediate goods firms rises modestly. The elasticity of substitution within sectors ϵ falls. I explore these results in more detail in Section 5.

	1990s	2000s	Meaning
ρ	**0.0147**	**0.001**	Rate of time preference (annual)
ϵ	**4.203**	**3.412**	Elasticity of substitution within sectors
η	**0.628**	**0.727**	Marginal cost of intermediate producers
δ_e	**0.096**	**0.087**	Exogenous entry/exit rate (annual)
λ	**1.053**	**1.064**	Min. qual. improvement
α	**4.826**	**3.745**	R&D cost parameter
ϕ	**0.789**	**1.692**	Shape of patent quality distribution

Table 4: Comparison of estimated parameters, 1990s vs. 2000s model steady states. Estimation details in the text and in Appendix C.2.

4.4 Model Validation

The model performs well in matching not just the average level of concentration but the entire (non-targeted) distribution of leaders' market shares across sectors of the economy in both periods. Figure 7 compares the distribution of leader market shares in the two study periods in the data and in the model. The shift is mainly due to increased average quality differences between leaders and followers in steady state, consistent with the findings of Andrews, Criscuolo, and Gal (2015) and Andrews, Criscuolo, and Gal (2016) that productivity differences within industries have grown. They also find that this divergence is particularly pronounced in ICT intensive sectors,

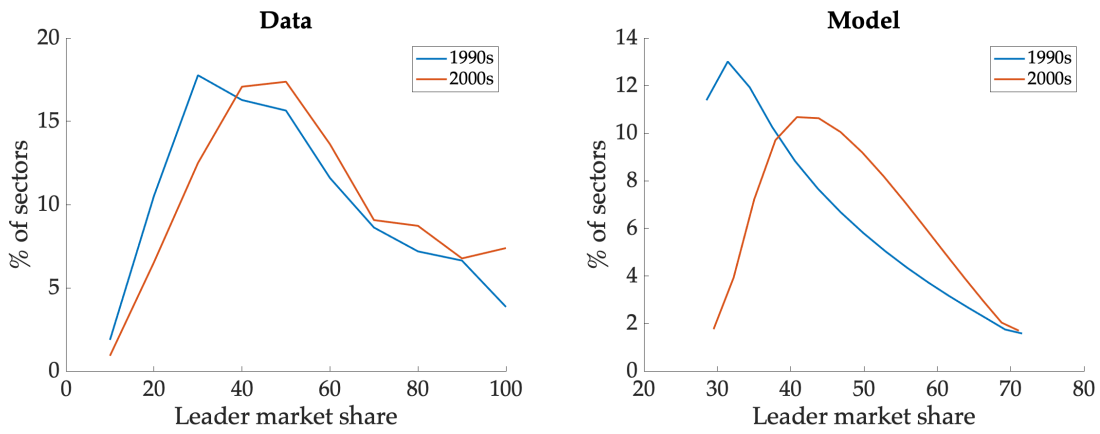

Figure 7: Shift in concentration across sectors, data vs. model. Leader market share is for 4-digit SIC industries in Compustat. See Appendix A.1 for a description of the sample.

and that sectors with wider productivity gaps have experienced deeper productivity slowdowns. Figure A.7 shows that rising concentration and productivity slowdowns are correlated at the sector level.

The model predicts that the average leader's share of total industry R&D rises from 17% to 64%. Among public firms, market leaders now perform a larger share of industry R&D than in the 1990s (Figure A.8), though the increase from 38% in 1999 to 50% in 2010 is smaller than in the model. Anderson and Kindlon (2019) also find a decline in R&D intensity among companies with fewer than 250 employees and an increase among larger firms in the National Science Foundation's Business R&D and Innovation Survey covering both public and private firms. Akcigit and Ates (2019) also document increasing concentration of patents among the top 1% of patenting firms and increasing flows of R&D employees from small to large firms.

5 Results

5.1 Decomposition

To understand the contribution of each estimated parameter change to matching the trends in the data, Table 9 in Appendix C.3 reports the effect of changing each parameter from its 1990s value to its 2000s value, holding the other parameters fixed at 1990s values. Note, however, that these are not the marginal effects of each parameter

on each moment; the moments are endogenously determined in steady state.

The decline in the discount rate ρ and the entry rate δ_e play similar roles in increasing incumbents' R&D expenditures to match the rise in R&D as a share of GDP, since all incumbents discount future profits less, increasing incentives for innovation. The decrease in the cost of R&D α also helps match the rise in R&D by increasing desired arrival rates of new ideas. Because they result in more R&D, these changes all have the effect of *raising* the TFP growth rate absent the other parameter changes. They do not substantially change the average level of concentration. Only the estimated changes in ϕ, governing relative patent quality of leaders and laggards, moves both concentration and productivity growth in the same direction as in the data.

5.2 Role of Changing Patent Quality

Table 5 summarizes the role of the model-implied change in ϕ compared to changes in all the other parameters at once to match the moments of interest. To decompose the effect of a change in the patent quality distribution, I compute the share of the changes in the data that are explained by a change in ϕ as follows:

$$\frac{M_j(\boldsymbol{\theta}_{1990s}, \phi_{2000s}) - M_j(\boldsymbol{\theta}_{1990s}, \phi_{1990s})}{D_{j,2000s} - D_{j,1990s}} \times 100, \tag{12}$$

where M_j is moment j in the model steady state with the other parameters $\boldsymbol{\theta}$ held fixed at their estimated 1990s values and $D_{j,t}$ denotes the moment's value in the data at time $t \in \{1990s, 2000s\}$.

A change in the patent quality distribution alone explains 145% of the productivity slowdown and 116% of the rise in concentration in the data. Turning to the mechanisms, laggard firms respond by choosing a lower innovation rate (Figure 8a).[22] The effect on leaders' innovation policies is ambiguous ex ante. Leaders are less likely to be overtaken so face less need to escape competition through innovation. Leaders with small technology gaps also face lower expected patent quality. Both of these forces lower desired innovation rates. On the other hand, facing a lower probability of being overtaken, leaders discount future gains from innovation less. In the calibrated model the latter force dominates and leaders choose a higher rate of innovations.

[22]Figure B.3 shows the impact of the estimated change in ϕ on expected quality improvements holding other parameters fixed.

Moment	Data	Decomposition	
	change (pp)	Innov. (%)	All others (%)
TFP growth	-0.9	145.5	-202.5
Leader market share	4.9	116.0	1.8
R&D/GDP	0.1	-1506.0	2960.8
Profits/GDP	1.4	54.7	68.0
R&D/sales	0.3	-1193.8	2327.4
Pat. qual., avg.	-11.7	109.5	-50.9
Pat. qual., followers	-12.6	171.7	-77.1
Pat. qual., leaders	-4.8	87.3	-48.1
Leadership turnover	-4.5	89.4	-43.9

Table 5: Share of changes in each data moment between 1990s and 2000s explained by the estimated parameter changes in Table 4. Column labelled "Innov." holds other parameters fixed at 1990s values while ϕ changes to its 2000s value. "All others" holds ϕ fixed and varies all other estimated parameters to 2000s values. Positive sign in the second and third columns indicates same direction of change as in the data. See text for more details.

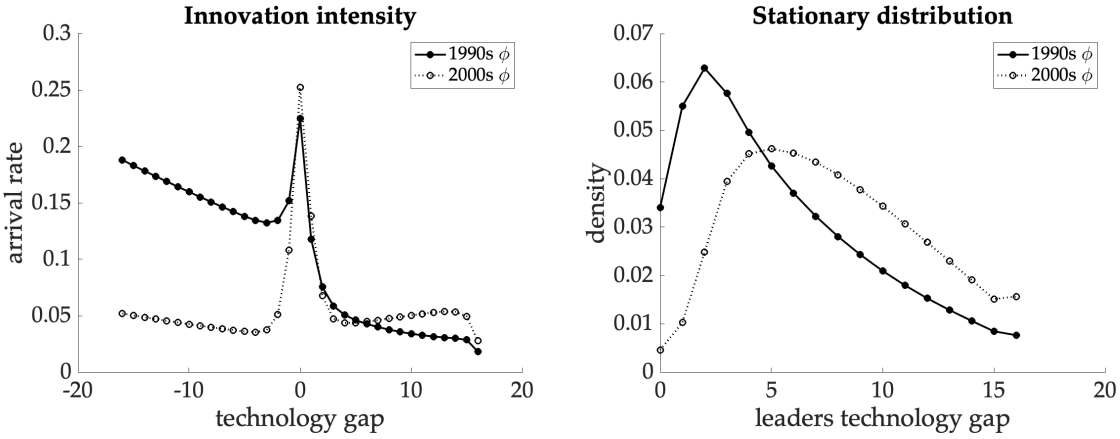

(a) Innovation policies **(b)** Stationary distribution

Figure 8: Innovation policies x as a function of the technology gap for all firms (panel a) and the distribution of sectors over the technology gap of the leader (panel b) comparing the 1990s steady state (solid line) in Table 1 to a counterfactual equilibrium where other parameters are fixed at 1990s values and ϕ changes to its 2000s value (dotted line).

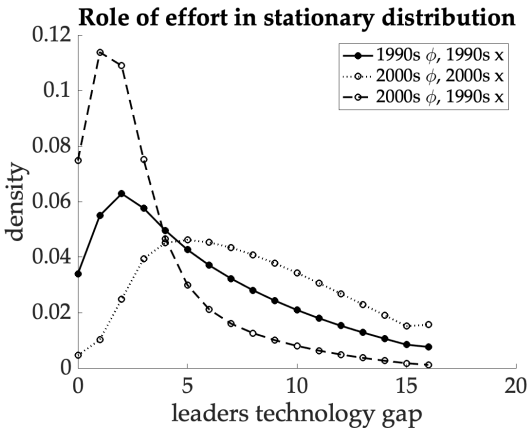

Figure 9: Distribution of sectors over the leader's technology gap. Solid line is the 1990s steady state (Table 1). Dotted line is a counterfactual equilibrium where ϕ changes to its 2000s value and other parameters are fixed. Dashed line is a counterfactual where innovation policies x are fixed at 1990s values and ϕ changes to its 2000s value (dashed line).

Together, these responses cause the stationary distribution to shift right: more sectors feature a leader that has high relative product quality (Figure 8b). The rise in concentration, average markups (from 23% to 26%), and the profit share in the equilibrium with lower patent quality is driven purely by this composition effect. Note that fixing innovation effort at its 1990s level in the model but increasing ϕ would result in a *tighter* distribution around the neck-and-neck state, because competitors pull away from each other less frequently under a higher ϕ (Figure 9).

The growth rate declines when laggards' patent quality is lower for two reasons. First, there is an endogenous effect that comes from changes in firms' innovation policies x (Figure 8a). R&D expenditures as a share of output decline from 1.8% to 0.3% (Table 5). The decline in R&D expenditures is concentrated among industry laggards, whose average R&D intensity declines from 6.5% in the 1990s to 0.6% in the counterfactual. Leaders' average R&D intensity declines from 2% to 1.2%. Leaders' improvements are more incremental on average, so both the level effect of reduced R&D and the reallocation effect contribute to slower productivity growth.

Second, lowering expected patent quality exogenously lowers the growth rate. To decompose the importance of these two channels, I use two decompositions. First, to isolate the strategic effect, I solve the model holding ϕ fixed but varying firms' innovation policies x to their values in the 2000s counterfactual. This explains 71.6%

of the total decline in productivity growth due to changing patent quality in the model. The other decomposition isolates the first order effect of reducing laggards' patent quality by fixing firm innovation policies at their 1990s values and reducing patent quality exogenously by changing ϕ to its 2000s value. This accounts for 76.4% of the productivity slowdown in the model. The fact that the two sum to more than 100% of the slowdown is due to interactions between the two in steady state.

5.3 Transition Dynamics for ϕ

To understand the dynamics of the transition from the high to low laggard patent quality equilibrium, I consider a permanent decrease in laggards' patent quality consistent with the increase in ϕ estimated in section 5 and with the pattern of declining patent quality in Figure 2 in section 2.2. I assume the transition takes 40 years, with ϕ increasing smoothly for the first two and a half years and then remaining permanently higher. After the initial surprise, firms anticipate the path for ϕ over the transition. Computational details for the transition analysis are presented in Appendix C.4.

Figure 10 shows the evolution of productivity growth and concentration. Productivity growth closely tracks the decline in patent quality, while concentration rises slowly and after 40 years has not yet reached its new steady state value. Innovations are somewhat infrequent, so it takes time for leaders to pull ahead and laggards to fall as far behind as they are in the new steady state on average. These patterns are consistent with Figure 1 which shows the TFP growth rate declining quickly in the early 2000s and then remaining roughly flat while concentration continued to rise through 2017. There is an initial boom in productivity growth similar to the 1990s data that is due to a boom in R&D investment when the transition begins. The boom occurs because firms have perfect foresight that patent quality is declining and that the terminal value of market leadership is higher than in the initial steady state.

5.4 Increasing Market Power or Superstar Firms?

Recent research has focused on the potential costs of rising market power for growth and welfare (de Loecker, Eeckhout, and Unger (2020); De Loecker, Eeckhout, and Mongey (2021); Eggertsson, Robbins, and Wold (2021); Edmond, Midrigan, and Xu

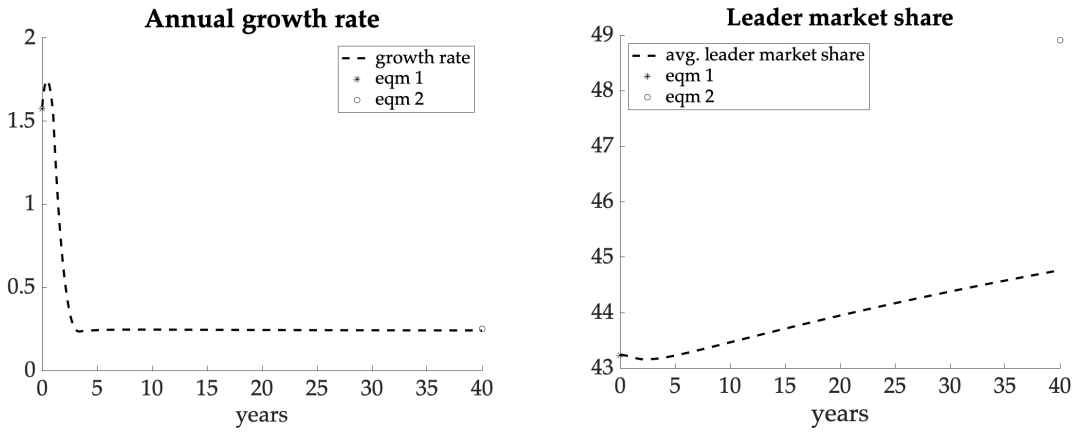

Figure 10: Transition dynamics for productivity growth and concentration over the transition to a lower level of patent quality. Details in the text.

(2021)). I model an increase in market power as a decrease in the substitutability of products in the same sector, ϵ, making the incumbents' varieties more differentiated and increasing the markup the leader charges for the same level of quality differences.

On the other hand, Autor et al. (2017) model the rise of superstar firms as an *increase* in product substitutability. They propose a static model where firms draw productivities prior to entry.[23] Firms that decide to enter produce differentiated goods and compete a la Bertrand. An increase in substitutability causes a reallocation of sales to the more productive "superstar" firms. Keeping the productivity distribution fixed, a sector's measured TFP rises unambiguously when substitutability increases for two reasons: the minimum productivity threshold for entrants rises and more productive firms increase their sales shares.

This static reallocation result holds in this model for the estimated 1990s parameter values when qualities are fixed. Increasing the substitutability of the varieties increases the leader's market share for a given level of the technology gap (Figure 11a) which raises sale-weighted and cost-weighted sector TFP.[24] On the other hand,

[23]The model in Autor et al. (2020) generalizes this formulation. The exercise presented here studies the specific shock to market toughness suggested in Autor et al. (2017). This shock could represent more fierce import competition from abroad, particularly China, in recent years. Amiti and Heise (2021) find support for this in Census data on U.S. and foreign manufacturing firms.

[24]Without the competitive fringe this is always true. But not necessarily with the fringe: if quality differences are small, increasing ϵ can cause a drop in the leader's market share.

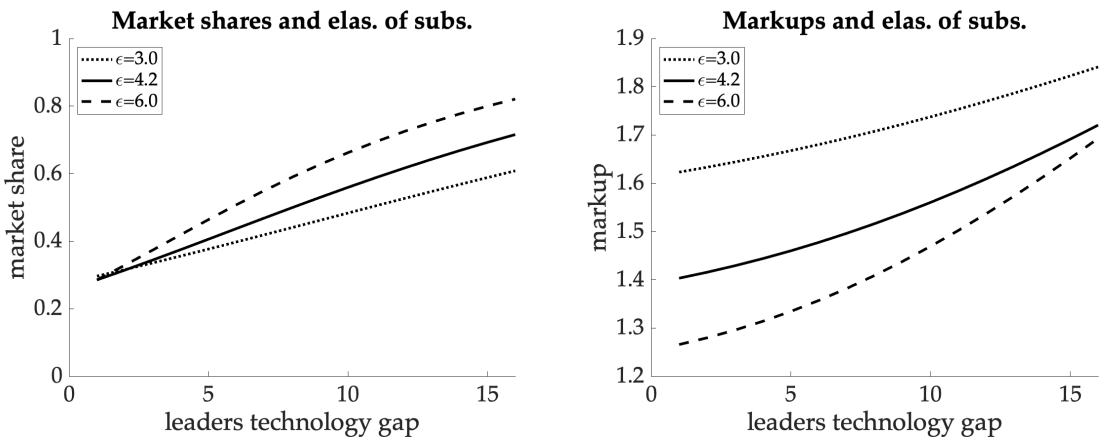

(a) Market shares **(b)** Markups

Figure 11: Static effect of the elasticity of substitution ϵ on market shares and markups of leaders using calibration in Table 1. $\epsilon = 3$ is the market power experiment (dotted line), $\epsilon = 4.2$ is the baseline (solid line), and $\epsilon = 6$ is the superstar firm experiment (dashed line).

decreasing ϵ (the "market power" experiment) reduces the leader's market share for a given level of the technology gap but increases their markup (Figure 11b).

Dynamically, however, firms will respond to changes in product differentiation and technology gaps will change, so the effects of these changes on concentration and growth on the BGP are ambiguous. Table 6 shows the effect of changes in ϵ on the model steady state. Under the market power experiment of decreasing ϵ, average markups rise by 10 percentage points, about a third of the total rise estimated by de Loecker, Eeckhout, and Unger (2020). Because leaders' markups are higher for the same level of the technology gap, this induces more innovation effort by laggard firms as they try to overtake the market leader (R&D/GDP rises from 1.8% to 2.7%, see Table 6 and Figure B.4a). This results in a faster growth rate, unlike in the data. There is also greater leadership turnover and average quality differences between leaders and followers go down (Figure B.4b), reducing market concentration.

The superstar firm shock, raising ϵ, lowers the growth rate and increases concentration, consistent with the data. The rise in concentration is due to two forces. First, the static reallocation force operates: even if technology gaps were unchanged, these same gaps would generate a higher average leader market share (Figure 11a). Second, changes in patenting frequency across different types of firms (Figure B.4a)

Moment	Model		
	$\epsilon = 3.0$	$\epsilon = 4.2$	$\epsilon = 6.0$
Avg. TFP growth, %	1.93	1.57	1.38
Avg. leader market share, %	39.67	43.23	49.34
R&D share of GDP, %	2.65	1.80	1.37
Profit share of GDP, %	6.31	5.69	5.61
Avg. R&D/sales, %	6.28	4.63	4.07
Avg. patent stock growth per patent, %	21.52	21.68	21.92
Avg. patent stock growth per patent, followers, %	33.33	33.72	34.36
Avg. patent stock growth per patent, leaders, %	10.07	9.99	9.86
Avg. leadership turnover, %	14.91	13.93	13.33
Avg. markup	1.33	1.23	1.18

Table 6: Steady state comparison, role of elasticity of subsitition ϵ, holding all other parameters fixed at the 1990s values in Table 1. $\epsilon = 3$ is the market power experiment, $\epsilon = 4.2$ is the baseline, and $\epsilon = 6$ is the superstar firm experiment.

cause the average technology gap to grow (Figure B.4b).

This latter effect occurs because the markup the leader charges is lower at all possible values of the technology gap when ϵ rises (Figure 11b), reducing the post-innovation gains to attaining market leadership and reducing the innovation effort of laggard firms (Figure B.4a). On the other hand, markups and profits become more elastic in the technology gap when ϵ is higher, and the likelihood of being overtaken falls, so leaders choose a higher arrival rate. There is both a reduction in the level and a shift in the location of R&D expenditures that generates a productivity slowdown, though this slowdown is much smaller than the slowdown driven by changing patent quality. Moreover, the average markup falls, contrary to the data.

6 Conclusion

This paper documents a decline in the contribution of new patents to firms' patent portfolios since 2000, suggesting that the quality of new ideas may have declined. Smaller firms drove the boom in innovation quality that has been attributed to the

arrival of information technology as a general purpose technology, as well as the bust that began in the early 2000s. This finding contributes to the debate on whether ideas are getting harder to find, emphasizing heterogeneity in this phenomenon across firms. Further work should investigate heterogeneity in the complementarity of general purpose technologies with firms' R&D investments, and determine whether similar patterns were present in the wake of previous general purpose technology waves.

To understand the consequences of this empirical fact I develop a general equilibrium growth model of innovations where multiple firms are active in each sector in each period. The quantitative model estimated for the U.S. in two different periods, the 1990s and the 2000s, points to declining patent quality as the main driver of rising concentration and the productivity slowdown over this period. Rising concentration in the 2000s is driven by a decline in the research effort of laggard firms and an increase by large firms, which causes average product quality differences between competitors to grow, consistent with the documented rise of superstar firms. Because leaders make smaller improvements on average, the economy grows more slowly.

I use the model to unify the endogenous growth literature with the growing literature on the rise of superstar firms. A rise in the elasticity of substitution has the potential to explain rising concentration and the productivity slowdown, providing a dynamic complement to the experiment in Autor et al. (2017), one that rationalizes the emergence of superstar firms alongside a productivity slowdown for a standard Schumpeterian reason: laggard firms' incentives to innovate fall because the value of market leadership is lower. However, this is inconsistent with patterns of markups and the profit share, which have risen rather than fallen over this period.

Analyzing welfare and optimal policy is left to future research, though the preceding discussion offers some insight into the relevant trade-offs between reducing static markup distortions and providing dynamic incentives to innovate. Such analysis should ensure that knowledge spillovers between firms and transition dynamics are properly accounted for. Another area for future research is how policy can spur the development of new general purpose technologies. Past work suggests this is difficult because there are significant positive externalities for other sectors that the inventor of the general purpose technology does not internalize. This paper suggests there may also be winners and losers within other industries, further complicating this problem.

References

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. "Innovation, Reallocation, and Growth." *American Economic Review* 108 (11): 3450–3491.

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

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt. 2014. "What Do We Learn From Schumpeterian Growth Theory?" Chapter 1 of *Handbook of Economic Growth*, Volume 2, 515–563. Elsevier B.V.

Aghion, Philippe, Stefan Bechtold, Lea Cassar, and Holger Herz. 2018. "The Causal Effects of Competition on Innovation: Experimental Evidence." *Journal of Law, Economics, and Organization* 34 (2): 162–195.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li. 2019a. "A Theory of Falling Growth and Rising Rents." *Federal Reserve Bank of San Francisco, Working Paper Series*, pp. 01–43.

Aghion, Philippe, Antonin Bergeaud, Mattheiu Lequien, and Marc J. Melitz. 2019b. "The Heterogenous Impact of Market Size on Innovation: Evidence from French Firm-Level Exports." *NBER Working Paper*.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. 2005. "Competition and Innovation: An Inverted-U Relationship." *Quarterly Journal of Economics* 120 (2): 701–728.

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

Akcigit, Ufuk, and Sina Ates. 2019. "What Happened to U.S. Business Dynamism?" *NBER Working Paper*, no. 25756.

———. 2020. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." *AEJ: Macroeconomics, Forthcoming*.

Akcigit, Ufuk, Sina Ates, and Giammario Impullitti. 2018. "Innovation and Trade Policy in a Globalized World." *NBER Working Paper*.

Akcigit, Ufuk, and William R. Kerr. 2018. "Growth Through Heterogeneous Innovations." *Journal of Political Economy* 126, no. 4.

Amiti, Mary, and Sebastian Heise. 2021. "U.S. Market Concentration and Import Competition." *Working Paper*.

Anderson, Gary, and Audrey Kindlon. 2019. "Indicators of R&D Performance by Size of Company." *National Science Foundation*, pp. 1–9.

Andrews, Dan, Chiara Criscuolo, and Peter Gal. 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, Dan, Chiara Criscuolo, and Peter N Gal. 2015. "Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence From OECD Countries." *The Future of Productivity: Main Background Papers*, pp. 1–38.

Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez. 2019. "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence." *AEJ: Macroeconomics* 11 (3): 67–110.

Arrow, Kenneth J. 1962. "The Economic Implications of Learning by Doing." *The Review of Economic Studies* 29 (3): 155–173.

Aum, Sangmin, Sang Yoon (Tim) Lee, and Yongseok Shin. 2018. "Computerizing industries and routinizing jobs: Explaining trends in aggregate productivity." *Journal of Monetary Economics* 97:1–21.

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 Firms." *Quarterly Journal of Economics* 135 (2): 645–709.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John van Reenen. 2017. "The Fall of the Labor Share and the Rise of Superstar Firms." *NBER Working Paper Series*, no. 23396.

Baqaee, David Rezza, and Emmanuel Farhi. 2020. "Productivity and Misallocation in General Equilibrium." *Quarterly Journal of Economics* 135 (1): 105–163.

Barkai, Simcha, and Seth G Benzell. 2018. "70 Years of US Corporate Profits." *Stigler Center New Working Paper Series*, no. 22.

Bartelsman, Eric J., and Mark E. Doms. 2000. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature* XXXVIII, no. September.

Berger, David W., Kyle F. Herkenhoff, and Simon Mongey. 2021. "Labor Market Power." *BFI Working Paper Series*.

Bessen, James. 2020. "Information Technology and Industry Concentration." *Journal of Law & Economics*, vol. 63.

Bessen, James, Erich Denk, Joowon Kim, and Cesare Righi. 2020. "Declining Industrial Disruption." *Boston University School of Law, Law & Economics Paper* 20, no. 28.

Bloom, Nicholas, Charles Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review* 110 (4): 1104–1144.

Bornstein, Gideon. 2021. "Entry and Profits in an Aging Economy: The Role of Consumer Inertia."

Bresnahan, Timothy F., and M. Trajtenberg. 1995. "General Purpose Technologies 'Engines of growth'?" *Journal of Econometrics* 65 (1): 83–108.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *AEJ: Macroeconomics*.

Carley, Michael, Deepak Hegde, and Alan Marco. 2015. "What is the Probability of Receiving a U.S. Patent?" *Yale Journal of Law and Technology*, vol. 203.

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

Cavenaile, Laurent, Pau Roldan, and Tom Schmitz. 2021. "International Trade and Innovation Dynamics with Endogenous Markups."

Chatterjee, Satyajit, and Burcu Eyigungor. 2020. "The Firm Size-Leverage Relationship and Its Implications for Entry and Business Concentration."

Comin, Diego, Javier Quintana, Tom Schmitz, and Antonella Trigari. 2021. "Measuring TFP: The Role of Profits, Adjustment Costs, and Capacity Utilization."

Corhay, Alexandre, Howard Kung, and Lukas Schmid. 2020. "Q: Risk, Rents, or Growth?"

Council of Economic Advisers. 2016. "Benefits of Competition and Indicators." *Issue Brief*, no. May.

Covarrubias, Matias, German Gutierrez, and Thomas Philippon. 2019. "From Good to Bad Concentration? U.S. Industries over the Past 30 Years." *NBER Macro Annual*.

Crouzet, Nicolas, and Janice Eberly. 2018. "Intangibles, Investment, and Efficiency." *AEA Papers and Proceedings* 108:426–431.

De Loecker, Jan, Jan Eeckhout, and Simon Mongey. 2021. "Quantifying Market Power and Business Dynamism." *Becker Friedman Institute Working Paper*, no. 2021-54.

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 U.S." *European Economic Review* 86:4–23.

de Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics* 135 (2): 561–644.

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

de Ridder, Maarten. 2021. "Market Power and Innovation in the Intangible Economy."

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2021. "How Costly Are Markups?"

Eggertsson, Gauti, Neil Mehrotra, and Jacob Robbins. 2019. "A Model of Secular Stagnation." *American Economic Journal: Macroeconomics* 11, no. 1.

Eggertsson, Gauti, Jacob Robbins, and Ella Getz Wold. 2021. "Kaldor and Piketty's Facts: The Rise of Monopoly Power in the United States." *Journal of Monetary Economics*.

Engbom, Niklas. "Firm and Worker Dynamics in an Aging Labor Market."

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

Ghazi, Soroush. 2021. "Large Firms and Long-Run Growth: How Trends in Market Structure Affect U.S. High-R&D Industries."

Goldberg, Jonathan, David Lopez-Salido, and Craig A Chikis. 2021. "Interest Rates and the Innovation Multiplier."

Graham, Stuart, and Deepak Hegde. 2015. "U.S. Patents and Their Disclosure Status." *Science* 347, no. 6219.

Griliches, Zvi. 1998. "Patent Statistics as Economic Indicators: A Survey." Chapter 13 of *R&D and Productivity: The Econometric Evidence*, 287–343. Univ. of Chicago Press.

———. 2001. *R&D, Education, and Productivity: A Retrospective*. Cambridge, Mass.: Harvard University Press.

Grullon, Gustavo, Yelena Larkin, and Roni Michaely. 2019. "Are US Industries Becoming More Concentrated?" *Review of Finance* 23 (4): 697–743.

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

Harris, Christopher, and John Vickers. 1987. "Racing with Uncertainty." *Review of Economic Studies* 54 (1): 1–21.

Helpman, Elhanan. 1998. *General Purpose Technologies and Economic Growth*. Cambridge: MIT Press.

Hoberg, Gerard, and Gordon Phillips. 2016. "Text-based network industries and endogenous product differentiation." *Journal of Political Economy* 124 (5): 1423–1465.

Hobjin, Bart, and Fernanda Nechio. 2019. "Sticker Shocks: Using VAT Changes to Estimate Upper-Level Elasticities of Substitution." *Journal of the European Economic Association* 17 (3): 799–833.

Hopenhayn, Hugo, Julian Neira, and Rish Singhania. 2018. "From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship and the Labor Share." *National Bureau of Economic Research Working Paper Series*, pp. 1–55.

Hsieh, Chang-Tai, and Esteban Rossi-Hansberg. 2021. "The Industrial Revolution in Services." pp. 139–158.

Jones, Callum, and Thomas Philippon. 2016. "The Secular Stagnation of Investment?"

Jones, Charles I. 2020. "The End of Economic Growth? Unintended Consequences of a Declining Population."

Jovanovic, Boyan, and Peter L Rousseau. 2005. *General Purpose Technologies*. Volume 1. Edited by Philippe Aghion and Steven N. Durlauf. Elsevier B.V.

Karahan, Fatih, Benjamin W. Pugsley, and Aysegul Sahin. 2019. "Demographic Origins of the Startup Deficit." *NBER Working Paper*, no. 25874.

Kehrig, Matthias, and Nicolas Vincent. 2021. "The Micro-Level Anatomy of the Decline of the Labor Share." *Quarterly Journal of Economics* 136 (2): 1031–1087.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring Technological Innovation over the Long Run †." *AER: Insights* 3 (3): 303–320.

Kelly, Bryan T., Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2018. "Measuring Technological Innovation over the Long Run." *SSRN Electronic Journal*.

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

Knott, Anne Marie, and Carl Vieregger. 2016. "Reconciling the Firm Size and Innovation Puzzle." *CES Working Paper*.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *Quarterly Journal of Economics*.

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

Liu, Ernest, and Song Ma. 2021. "Innovation Networks and Innovation Policy."

Liu, Ernest, Atif R. Mian, and Amir Sufi. 2021. "Low Interest Rates, Market Power, and Productivity Growth." *Econometrica*.

Luttmer, Erzo G.J. 2007. "Selection, Growth, and the Size Distribution of Firms." *Quarterly Journal of Economics*, no. August:1103–1145.

Mongey, Simon. 2021. "Market Structure and Monetary Non-neutrality."

Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–1297.

Pellegrino, Bruno. 2021. "Product Differentiation, Oligopoly, and Resource Allocation."

Peters, Michael. 2020. "Heterogeneous Mark-ups, Growth and Endogenous Misallocation." *Econometrica* 88, no. 5.

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

Peters, Ryan H., and Lucian A. Taylor. 2017. "Intangible Capital and the Investment-q Relation." *Journal of Financial Economics* 123 (2): 251–272.

Pugsley, Benjamin W., Petr Sedlacek, and Vincent Sterk. 2021. "The Nature of Firm Growth." *American Economic Review* 111 (2): 547–579.

Rosen, Richard J. 1991. "Research and Development with Asymmetric Firm Sizes." *RAND Journal of Economics* 22 (3): 411–429.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter. 2020. "Diverging Trends in National and Local Concentration." *NBER Macroeconomics Annual*, vol. 35.

Syverson, Chad. 2017. "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown." *Journal of Economic Perspectives*.

Zachariadis, Marios. 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.

Zhang, Linyi. 2018. "Escaping Chinese Import Competition? Evidence from U.S. Firm Innovation." *Working Paper*, pp. 1–69.

A Online Data Appendix

A.1 Data Description

The main source of data for the paper is the Compustat Fundamentals Annual database, 1962-2017 (though most analysis focuses on the post-1980 period). I restrict attention to firms incorporated in the U.S. (FIC="USA") reporting in U.S. dollars (CURCD="USD"). I further drop financial (those with 4 digit SIC code beginning with 6 and those with INDFMT="FS"), agricultural (those with 4 digit SIC<1000), utilities firms (those with 4 digit SIC code beginning with 49) and non-operating establishments (those with 4 digit SIC code=9995). This data is merged with the updated patent data through 2020 from Kogan et al. (2017).[25]

A.2 Data Sources and Moment Computations

Table 7 lists the source and, where necessary, computation method for each target moment from the data.

[25]The updated data can be found at: https://github.com/KPSS2017/Technological-Innovation-Resource-Allocation-and-Growth-Extended-Data.

Moment	Source	Computation/Series Name
TFP growth	Comin et al. (2021)	Utilization-adjusted annual total factor productivity growth estimates allowing for non-zero profits
Leader market share	Compustat	Average of sales share (SALE) of largest firm in each 4-digit SIC industry (weighted by industry size in terms of sales)
Patent quality $\equiv$ patent stock growth per patent (psgpp)	Kogan et al. (2017)	$rTsm_{it} = \frac{Tsm_{it}}{GDPdefl_t}$ is the real value of firm i's patents issued in year t. $psgpp_{it} = \frac{\frac{rTsm_{it}}{fNpats_{it}}}{\sum_{s=1}^{t-1} rTsm_{is}}$; s =first year in Compustat. Citation-based version substitutes Tcw (not deflated). I winsorize the top and bottom 1% of observations of psgpp.
R&D share of GDP	OECD Main Science and Technology Indicators	Business Expense R&D (private)/GDP
R&D intensity	Compustat	XRD/SALE, winsorizing the top and bottom 10% of observations, mean for all firms, assuming 0 if XRD missing.
Profit share of GDP	Bureau of Economic Analysis/FRED	Profits after tax with inventory valuation and capital consumption adjustments/Gross domestic income
Leader's share of R&D	Compustat	Average sales leader share of total R&D in 4-digit sector (weighted by industry size in terms of sales)
Leadership turnover	Compustat	Share of 4-digit SIC industries with new sales leader per year

Table 7: Data sources and computation method for each moment used in the text.

A.3 Additional Patent Quality Figures

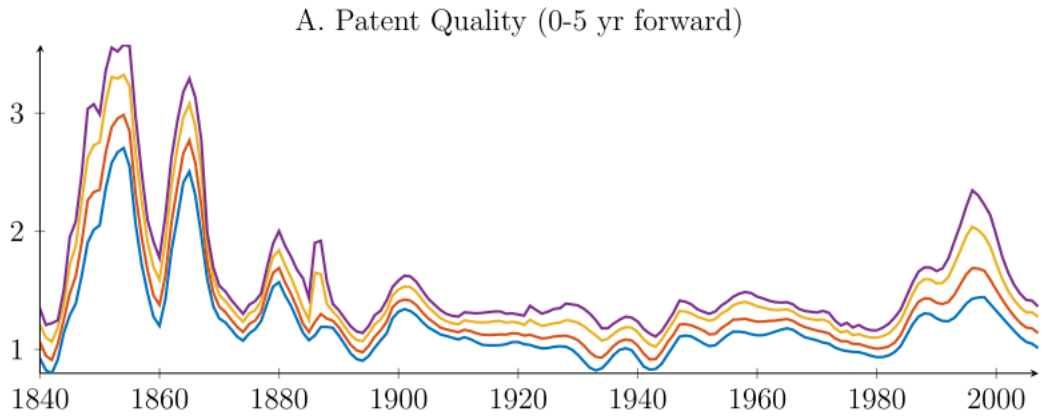

Figure A.1: Percentiles of text-based patent quality distribution over time. Blue = P50, Red = P75, Yellow = P90, Purple = P95. Source: Kelly et al. (2018) Figure 3a.

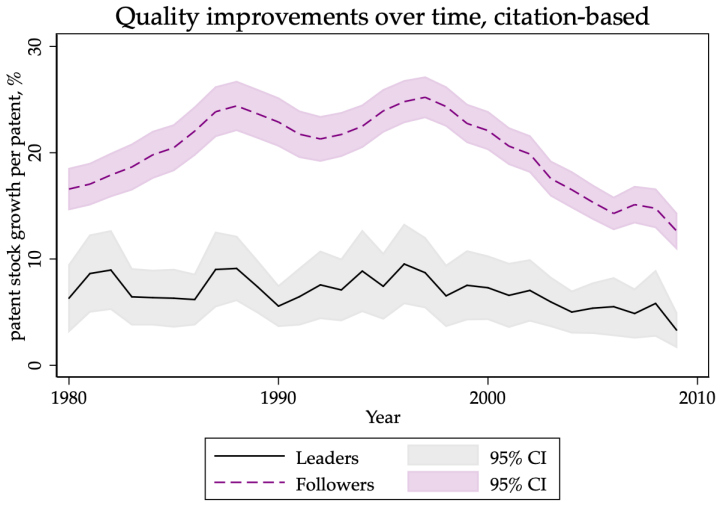

Figure A.2: Contribution of average new patent to firm's existing stock of patents, substituting forward citations counts (Tcw) for dollar value, Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries in a given year and followers are all other public firms in the sample. Three year moving averages. See Appendix A.1 for more details.

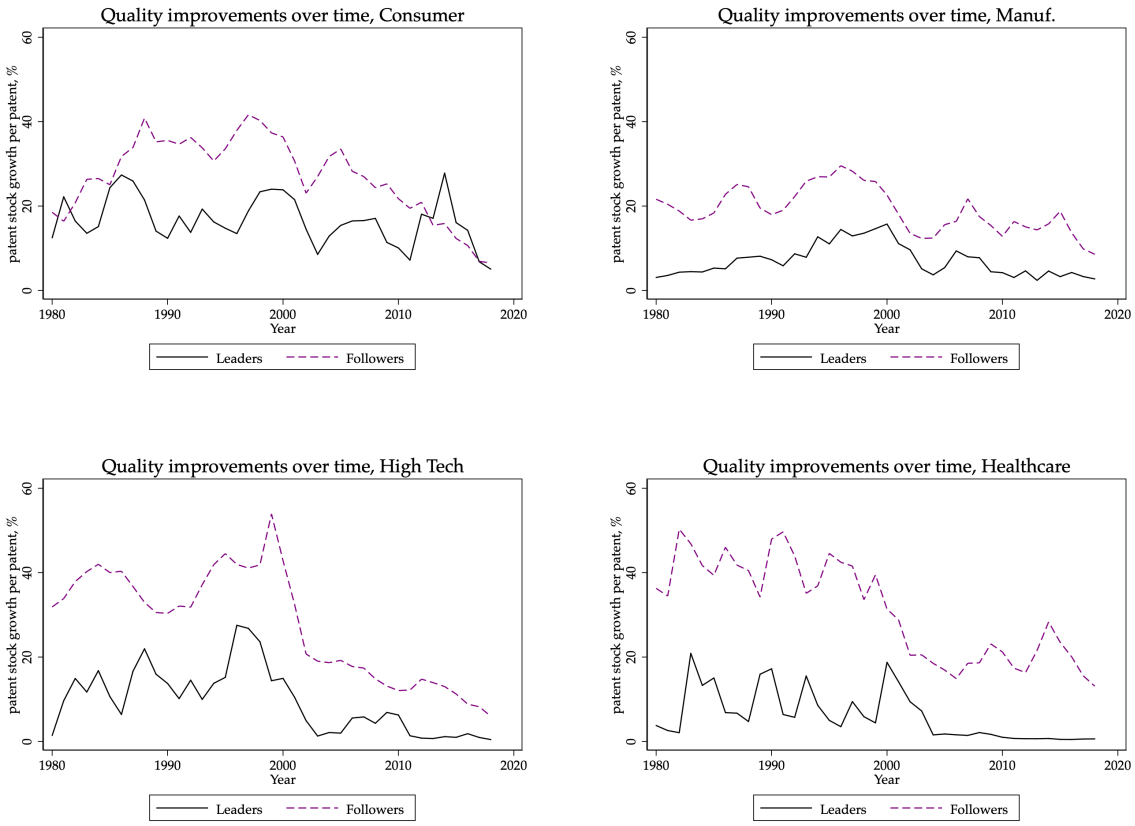

Figure A.3: Average patent quality differences between leaders and followers in Fama-French 5 broad industry categories (excluding "Other" category), using estimated patent values from Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries in a given year and followers are all other public firms in the sample. Three year moving averages. See Appendix A.1 for more details.

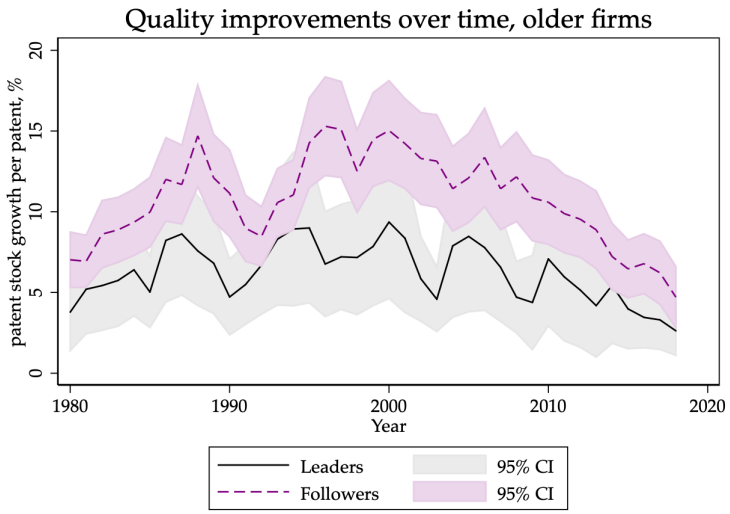

Figure A.4: Contribution of average new patent to value of firm's existing stock of patents, using estimated patent values from Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries in a given year and followers are all other public firms in the sample, restricting attention to firms that have been public at least 20 years in the year patent is issued. See Appendix A.1 for more details.

A.4 TFP and Markup Estimation

I use the Compustat data on U.S. public firms with the restrictions described above from 1962-2018 to estimate revenue-based total factor productivity (TFPR) and markups at the firm level. In the estimation I further exclude firms without an industry classification, firm-years where acquisitions are more than 5% of assets ($ACQ/AT > 0.05$), and firm-years with non-positive employment ($EMP \leq 0$) or negative sales ($SALE \leq 0$). The sample includes around 3,000 firms per year, though this number varies over time.

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

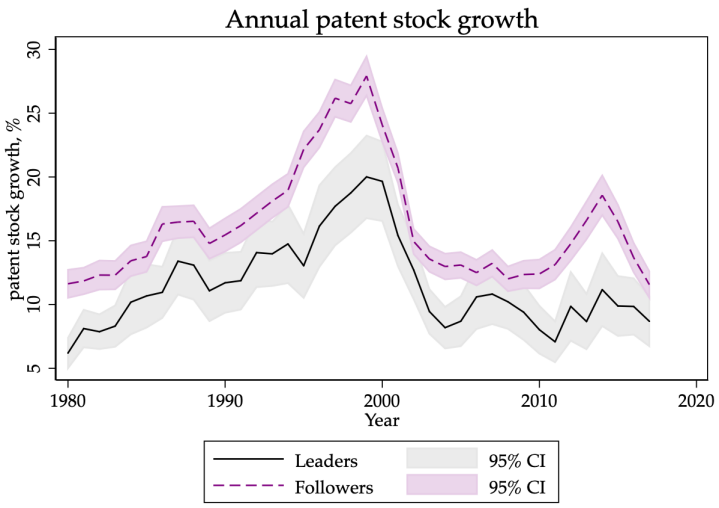

Figure A.5: Average annual growth of firm's patent stock conditional on patenting at least once in that year, using estimated patent values from Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries in a given year and followers are all other public firms in the sample. Three year moving averages. See Appendix A.1 for more details.

Economic Analysis (BEA) deflator for non-residential fixed investment.

In de Loecker and Warzynski (2012) the authors show that under a variety of pricing models firm i's markup at time t, μ_{it}, can be computed as a function of the output elasticity θ_{it}^V of any variable input and the variable input's cost share of revenue[26] :

$$\mu_{it} = \theta_{it}^V \frac{P_{it}Q_{it}}{P_t^V V_{it}} \tag{13}$$

where P_{it} is the output price of firm i's good at time t, Q_{it} its output, P_t^V the price of the variable input and V_{it} the amount of the input used.

Following de Loecker, Eeckhout, and Unger (2020) 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

[26]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 need be assumed.

would be needed to compute the labor cost share needed to compute the markup simultaneously with productivity.

For the results presented in this paper, I assume a Cobb-Douglas production function[27] for firm i in 2-digit SIC sector s in year t so that factor shares may vary across sectors but not over time:

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

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

$$y_{i,t} = \alpha + \eta_t + \delta_i + \beta_{M,s} m_{i,t} + \beta_{K,s} k_{i,t-1} + \varepsilon_{i,t}.$$

In the above equation, $\beta_{M,s}$ captures the sector specific variable output elasticity, so I use equation 13 to obtain the markup from the estimated $\hat{\beta}_{M,s}$ and the inverse cost share $\frac{SALE}{COGS}$.

A.5 Industry Profit Shares

The competitive fringe 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 A.6).[28]

A.6 Additional Model Validation Figures

An empirical exploration of the causal relationships among productivity growth and concentration is beyond the scope of this paper. However, given the sectoral heterogeneity in the decline in laggards patent quality in Figure A.3 suggesting that this phenomena differs across industries, we might expect rising concentration and the

[27] Alternative estimation of a translog production function yielded similar estimates.

[28] Following the estimation strategy of de Loecker, Eeckhout, and Unger (2020) to estimate TFP and markups at the firm level, TFP and sales share are correlated, and the figure looks similar if one uses a productivity ranking instead of sales-share based ranks.

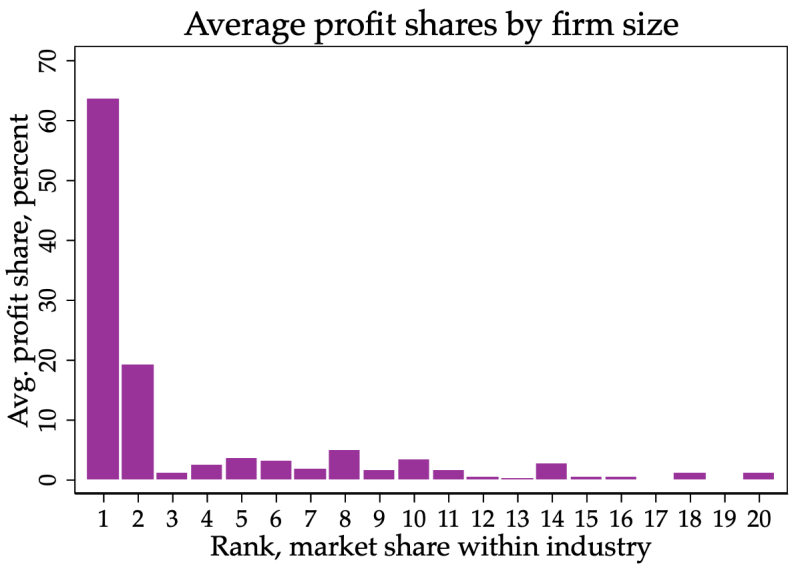

Figure A.6: Firms are ranked by market share (sales) within 4-digit SIC industries for a given year, and these ranks are compared to profit shares (firm's own operating income (OIDBP) as a share of industry-total operating income). Averages across 4-digit sectors and across time for 1980-2017. See Appendix A.1 for more details.

productivity slowdown to be correlated at the sector level. I use data from Bureau of Economic Analysis estimates of multifactor productivity[29] at the 3-digit NAICS level and Compustat to check the association between the change in the leader's market share and the change in the sector's average productivity growth rate from 1994-2003 to 2004-2017. Sectors with greater slowdowns in productivity growth also saw greater increases in concentration on average (Figure A.7).

[29]https://www.bea.gov/data/special-topics/integrated-industry-level-production-account-klems

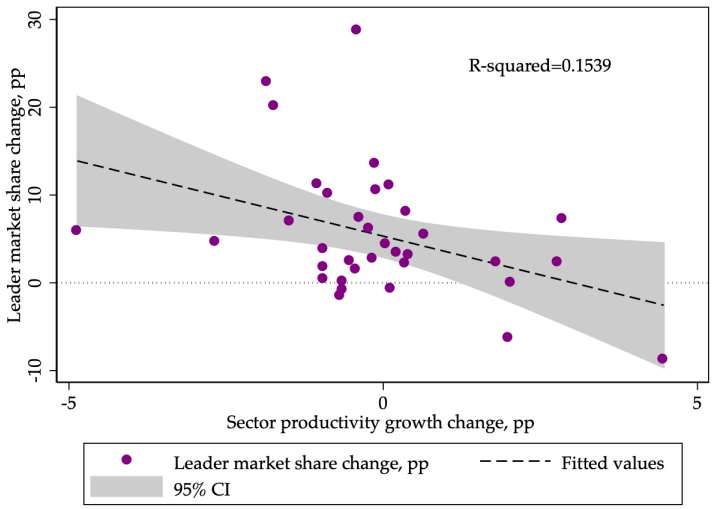

Figure A.7: Each dot is a 3-digit NAICS sector. Dots show the change in average annual total factor productivity growth from the BEA Integrated Industry-Level Production Accounts between 1994-2003 and 2004-2017 (x-axis) plotted against the change in the largest firm's average share of industry sales between 1994-2003 and 2004-2017 in Compustat (y-axis). Excludes oil, gas, and mining industries (211, 212, 213).

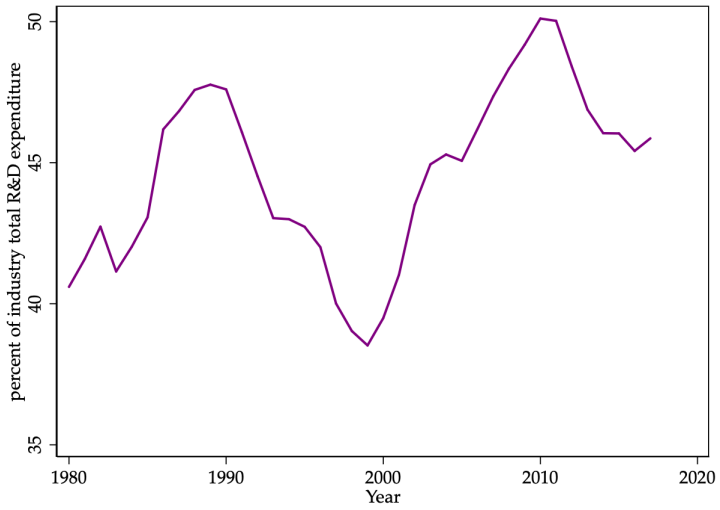

Figure A.8: Research and development expenditures (XRD) of sales leaders in 4-digit SIC industries in Compustat as a share of total R&D expenditures of all public firms in that sector. Average across industries, sale-weighted by industry size. See Appendix A.1 for more details.

B Online Model Appendix

B.1 Proof Prices Depend on Relative Quality

Relative quality refers to the ratio of qualities of the two incumbent firms in a sector (dropping the sector notation j) $\frac{q_1}{q_2}$ for firm 1 and $\frac{q_2}{q_1}$ for firm 2. Below I show that the firms' pricing strategies depend only on relative quality, not the level of their own or their rival's quality.

First, this is clearly satisfied for the technology follower ($m_i < 0$) who sets price equal to marginal cost η regardless of absolute quality, and for sectors where $m_1 = m_2 = 0$, that is, when firms are neck-and-neck, because of the presence of the competitive fringe.

For the leader ($m_i > 0$), plugging the final good firm's demand for good i into the definition of the market share and using the definition of the price index yields:

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

where $-i$ denotes the follower. Now using the pricing decision of the leader:

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

Thus there is a mapping from technology gaps to market shares and prices that is independent of quality levels. ■

B.2 Value Function Boundary Equations

For a firm that's the furthest possible number of steps behind its competitor (that is, at gap $-\bar{m}$) and absolute quality level q_{ijt}:

$$\begin{aligned} r_t V_{-\bar{m},t}(q_{ijt}) - \dot{V}_{-\bar{m},t}(q_{ijt}) = \max_{x_{-\bar{m},t}} \{ & 0 - \alpha \frac{(x_{-\bar{m},t})^\gamma}{\gamma} q_{ijt}^{\frac{1}{\beta}-1} \\ & + x_{-\bar{m},t} \sum_{n_t=-\bar{m}+1}^{\bar{m}} \mathbb{F}_m(n_t)[V_{nt}(\lambda^{n_t-(-\bar{m})} q_{ijt}) - V_{-\bar{m},t}(q_{ijt})] \\ & + x_{\bar{m},t}(V_{-\bar{m},t}(\lambda q_{ijt}) - V_{-\bar{m},t}(q_{ijt})) \\ & + \delta_e(0 - V_{-\bar{m},t}(q_{ijt})\}. \end{aligned}$$

The difference between this and equation 5 is in the third line, where if the firm's competitor innovates, there is a spillover that causes the firm at gap $-\bar{m}$ to improve its quality by λ.

For a firm at gap $\bar{m}$ the value function is:

$$\begin{aligned} r_t V_{\bar{m},t}(q_{ijt}) - \dot{V}_{\bar{m},t}(q_{ijt}) = \max_{x_{\bar{m},t}} \{ & \pi(\bar{m}, q_{ijt}) - \alpha \frac{(x_{\bar{m},t})^\gamma}{\gamma} q_{ijt}^{\frac{1}{\beta}-1} \\ & + x_{\bar{m},t}(V_{\bar{m},t}(\lambda q_{ijt}) - V_{\bar{m},t}(q_{ijt})) \\ & + x_{-\bar{m},t} \sum_{n_t=-\bar{m}+1}^{\bar{m}} \mathbb{F}_{-\bar{m}}(n_t)[V_{nt}(q_{ijt}) - V_{\bar{m},t}(q_{ijt})] \\ & + \delta_e(0 - V_{\bar{m},t}(q_{ijt})\}, \end{aligned}$$

where:

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

B.3 Derivation of Final Output

The section derives the expression for final output in equations 6 and 7. Dropping the time subscript t, plugging the pricing strategies in equation 4 and $p_i = \eta$ for firms with $m_i \leq 0$ into the demand curve (3) to obtain the output of each incumbent and

plugging these outputs into equation (2) and equation (2) into equation (1) simplifies as:

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

The demand shifter $P^{\frac{1}{\beta}}L$ index is common to all firms and can therefore be factored out of the integral (and normalized to one since I assume that there is no population growth). The quality-adjusted price index P_j of each sector falls as the qualities of the two products in the sector grow, and the exponent is negative for all $\beta \in (0,1)$ so Y is increasing in firms' qualities.

Common to all firms with a particular technology gap m are the prices $p(m)$ of the firm at gap m and its competitor at $-m$, $p(-m)$. At time t, therefore, Y can be expressed as:

$$Y_t = \frac{1}{2}\frac{L}{1-\beta} P^{\frac{1-\beta}{\beta}} \sum_{m=-\bar{m}}^{\bar{m}} \left(\int_0^1 \left(q_{it}^{\epsilon-1} p_i(m)^{1-\epsilon} + q_{-it}^{\epsilon-1} p_{-i}(-m)^{1-\epsilon}\right)^{-\frac{(1-\beta)}{\beta(1-\epsilon)}} \mathbb{1}_{\{i\in\mu_{mt}\}} di\right)$$

where μ_{mt} is the measure of firms at technology gap m at time t and the above integration is taken over firms rather than sectors. More simply:

$$Y_t \equiv \frac{1}{2}\frac{L}{1-\beta} P^{\frac{1-\beta}{\beta}} \sum_{m=-\bar{m}}^{\bar{m}} Q_{mt},$$

where Q_{mt} is defined as:

$$\begin{aligned}
Q_{m,t} &= \int_0^1 \left(q_{it}^{\epsilon-1} p(m)^{1-\epsilon} + q_{-it}^{\epsilon-1} p(-m)^{1-\epsilon}\right)^{-\frac{(1-\beta)}{\beta(1-\epsilon)}} \mathbb{1}_{\{i\in\mu_{mt}\}} di \\
&= \left(p(m)^{1-\epsilon} + (\lambda^{-m})^{\epsilon-1} p(-m)^{1-\epsilon}\right)^{\frac{1-\beta}{\beta(\epsilon-1)}} \int_0^1 q_{i,t}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i\in\mu_{mt}\}} di.
\end{aligned}$$

B.4 Output Growth on Balanced Growth Path

To understand how aggregate output evolves, this section studies the evolution of $\tilde{Q}_{m,t}$ (defined in equation 11) between t and $t+dt$ for all m. These expressions are similar to those for the stationary distribution (equations 8-10) because they are based on the movement of firms to different technology gaps from their rival, but account for the quality improvements that occur because of innovation.

Assuming fixed distribution $\mu_{mt} = \mu_m$ for all m, t:

$$\dot{\tilde{Q}}_{mt} = \int_0^1 q_{m,t+dt,i}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i\in\mu_m\}} di - \int_0^1 q_{m,t,i}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i\in\mu_m\}} di.$$

that is, quality growth at gap m is due to the change an index of the qualities of all the firms with technology gap m. Consider an arbitrary $m \in (-\bar{m}, \bar{m})$ ($-\bar{m}$ and $\bar{m}$ are special cases because of spillovers). A portion of firms at m at t innovate to a different gap, and another portion leave gap m because their competitor innovates. Because all firms at gap m choose the same arrival rate x_m, these are a random sample of the firms at gap m at time t. The outflows from $\dot{\tilde{Q}}_m$ are:

$$-(x_m + x_{-m}) \int_0^1 q_{m,t,i}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i\in\mu_m\}} di = -(x_m + x_{-m})\tilde{Q}_m.$$

The inflows to m's quality index come from two sources. First, some firms innovate into position m from a lower position n, improving their quality by λ^{m-n}. The probability they innovate and reach gap m is given by $x_n F_n(m)$. Some firms fall back to m from a higher gap n because their competitor innovates to $-m$. The probability their competitor reaches $-m$ is given by $x_{-n}F_{-n}(-m)$. So cumulative inflows are:

$$\sum_{n=-\bar{m}}^{m-1} x_n F_n(m)(\lambda^{(m-n)})^{\frac{1-\beta}{\beta}} \tilde{Q}_n + \sum_{n=m+1}^{\bar{m}} x_{-n}F_{-n}(-m)\tilde{Q}_n.$$

Putting it together:

$$\dot{\tilde{Q}}_{mt} = \sum_{n=-\bar{m}}^{m-1} x_n F_n(m)(\lambda^{(m-n)})^{\frac{1-\beta}{\beta}} \tilde{Q}_n + \sum_{n=m+1}^{\bar{m}} x_{-n}F_{-n}(-m)\tilde{Q}_n - (x_m + x_{-m})\tilde{Q}_m. \quad (14)$$

At the lowest gap there are spillovers when the competitor innovates:

$$\dot{\tilde{Q}}_{-\bar{m}t} = \sum_{n=-\bar{m}+1}^{\bar{m}} x_{-n}F_{-n}(\bar{m})\tilde{Q}_n + x_{\bar{m}}(\lambda^{\frac{1-\beta}{\beta}} - 1)\tilde{Q}_{-\bar{m}} - x_{-\bar{m}}\tilde{Q}_{-\bar{m}}. \quad (15)$$

At the highest gap the leading firm does not exit that gap when they innovate:

$$\dot{\tilde{Q}}_{\bar{m}t} = \sum_{n=-\bar{m}}^{\bar{m}-1} x_n F_n(\bar{m})(\lambda^{(m-n)})^{\frac{1-\beta}{\beta}} \tilde{Q}_n + x_{\bar{m}}(\lambda^{\frac{1-\beta}{\beta}} - 1)\tilde{Q}_{\bar{m}} - x_{-\bar{m}}\tilde{Q}_{\bar{m}}. \tag{16}$$

Given equations 14, 15, and 16, on a balanced growth path where $\frac{\dot{\tilde{Q}}_{mt}}{Y_t}$ is constant, it's sufficient to assume $\frac{\tilde{Q}_{mt}}{Y_t}$ is constant over time for all $m \in [-\bar{m}, \bar{m}]$. Differentiating $\frac{\tilde{Q}_{mt}}{Y_t}$ with respect to time yields:

$$\begin{aligned}(\frac{\dot{\tilde{Q}}_m}{Y}) &= \frac{\dot{\tilde{Q}}_m}{Y} - \frac{\tilde{Q}_m}{Y}\frac{\dot{Y}}{Y} \\ &= \frac{\dot{\tilde{Q}}_m}{Y} - g\frac{\tilde{Q}_m}{Y}.\end{aligned}$$

Imposing that the left hand side is zero implies:

$$\frac{\dot{\tilde{Q}}_m}{Y} = g\frac{\tilde{Q}_m}{Y}.$$

The vector on the left hand side is defined above by the flow equations (14), (15), and (16) divided by output. Use those equations to form a matrix A that captures the flow equations:

$$\frac{\dot{\tilde{Q}}_m}{Y} = A\frac{\tilde{Q}_m}{Y} = g\frac{\tilde{Q}_m}{Y}.$$

The values in A depend on λ, ϕ, and x_m. The above equation means that the growth rate g is an eigenvalue of the matrix A and $\frac{\tilde{Q}_m}{Y}$ is the corresponding eigenvector of A. If there is only one positive, real eigenvalue of A then there is exists only one such balanced growth path where the contribution of the growth of the quality index of each technology gap to the total growth rate is constant and the growth rate of the economy is constant.

B.5 Alternate Model With No Competitive Fringe

I solve the full dynamic model without the competitive fringe so that both firms exercise market power over their variety i of sector j's good. There is still the possibility of exogenous entry/exit, though this assumption can be relaxed as well. The analogy from the model to the data becomes less obvious under this assumption, since the

laggard firm can no longer be thought of representing many firms producing generic products that are perfectly substitutable with other generic products but imperfectly substitutable with the brand produced by the leader. In this setup the quality leader always has at least 50% market share, unlike in the data. This assumption also gives empirically counterfactual predictions that the profit shares of total industry profits of the market leader and the other firm in the industry are relatively similar, contradicting the pattern shown in Figure A.6.

Nonetheless, the main results carry through under this alternate assumption. Before describing these alternate results, I return to the pricing problem of the firms assuming the follower can now choose its optimal markup. Using the same derivation as in section 3.3.1 it can be shown that both firms follow the pricing policy the leader follows in the baseline model:

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

where

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

I look for a Markov perfect equilibrium with balanced growth where each firm's price is the best response to its competitor's price at time t. The algorithm for finding the steady state remains the same, plugging in the pricing functions of the firms, illustrated in Figure B.1.

Table 8 gives the results of the same experiment as in section 4.3 under the alternate pricing strategies with the same parameters as in Table 1 and Figure B.2 shows the policy functions and stationary distributions. Note that the escape competition motive around the neck and neck state disappears in the version without the competitive fringe. As before, changing ϕ has a level effect on total innovation effort but also changes the location of R&D from laggard firms to leading firms.

The level of the growth rates and the change in the growth rate from one steady state to the other due to a change in ϕ under Bertrand pricing are very similar to the baseline model with marginal cost pricing of the follower. The increase in concentration is smaller since the change in technology gaps is not as large as in the main case (Figure B.2), though technology gaps do increase modestly. As for

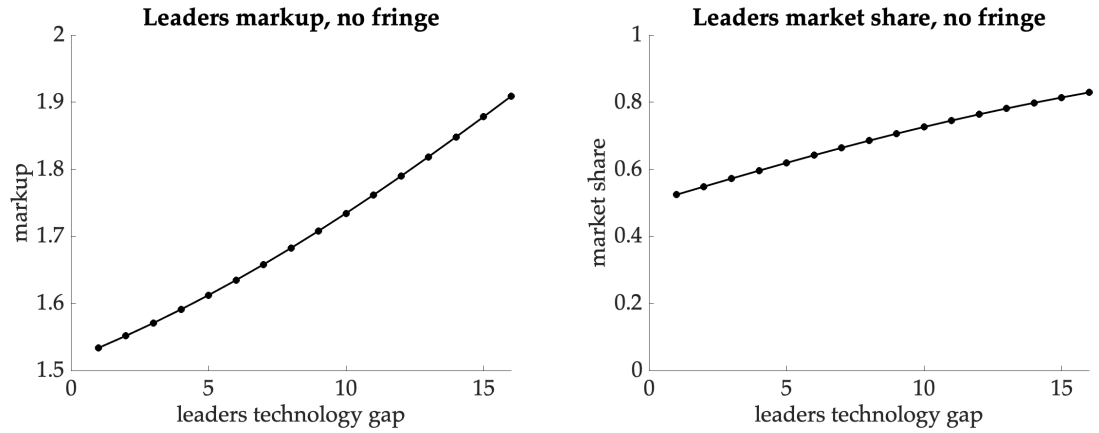

Figure B.1: Leader's markup (price over marginal cost) and market share as a function of the leader's technology gap m in the model using the 1990s calibration in Table 1 in the version of the model with no competitive fringe.

the growth decomposition, the effects of the firms' innovation responses is smaller (58.7%), and the first order effect of lowering the probability of radical innovations is a bit larger than in the baseline model with the competitive fringe (85.1%).

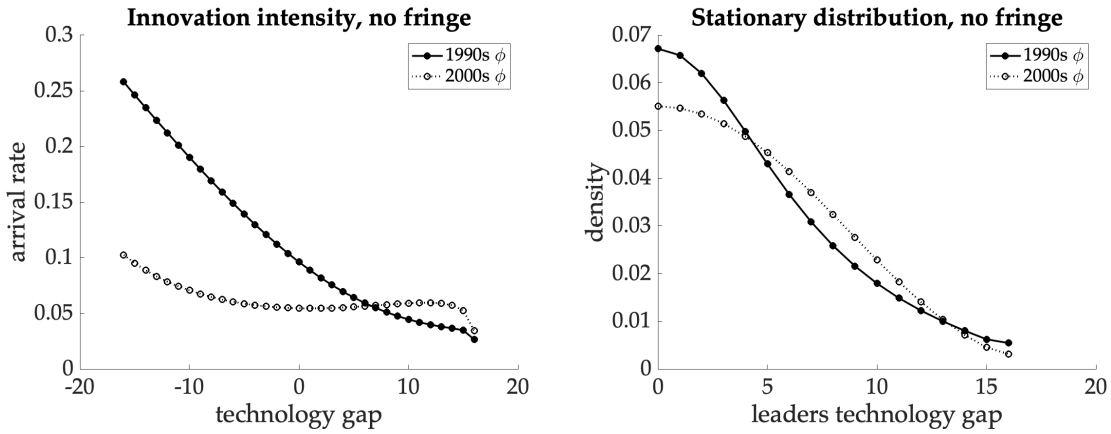

Figure B.2: Innovation policies x as a function of the technology gap for all firms (panel a) and the stationary distribution of sectors over the technology gap of the leader (panel b) using the 1990s calibration in Table 1 in the version of the model with no competitive fringe.

Moment	Data			Model		
	1990s	2000s	Chg. (pp)	1990s	2000s ϕ	Chg. (pp)
TFP growth	1.57	0.66	-0.91	1.59	0.30	-1.29
Leader market share	43.23	48.12	4.89	61.58	62.43	0.85
R&D/GDP	1.80	1.89	0.09	2.03	0.56	-1.47
Profits/GDP	6.61	6.61	1.37	14.47	14.40	-0.06
R&D/sales	3.75	4.06	0.30	6.99	1.62	-5.37
Pat. qual., avg.	22.90	11.22	-11.68	20.90	8.04	-12.86
Pat. qual., followers	24.86	12.23	-12.63	32.11	10.20	-21.91
Pat. qual., leaders	10.00	5.17	-4.84	10.30	6.00	-4.30
Leadership turnover	13.74	9.28	-4.47	13.98	10.92	-3.06

Table 8: "Data" columns show the levels and changes in the key moments in the data. "Model" columns show the levels and changes of the moments in the equilibrium of the model without the competitive fringe under the parameterization in Table 1 ("1990s") and the counterfactual equilibrium with the same parameterization except varying ϕ to its estimated value in the 2000s from Table 4 ("2000s ϕ").

B.6 Extra Model Figures

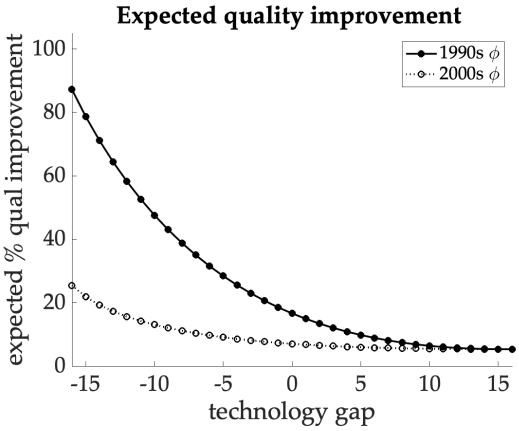

Figure B.3: Expected quality improvement from innovation as a function of the technology gap comparing the estimated values of ϕ, holding other parameters fixed at 1990s values.

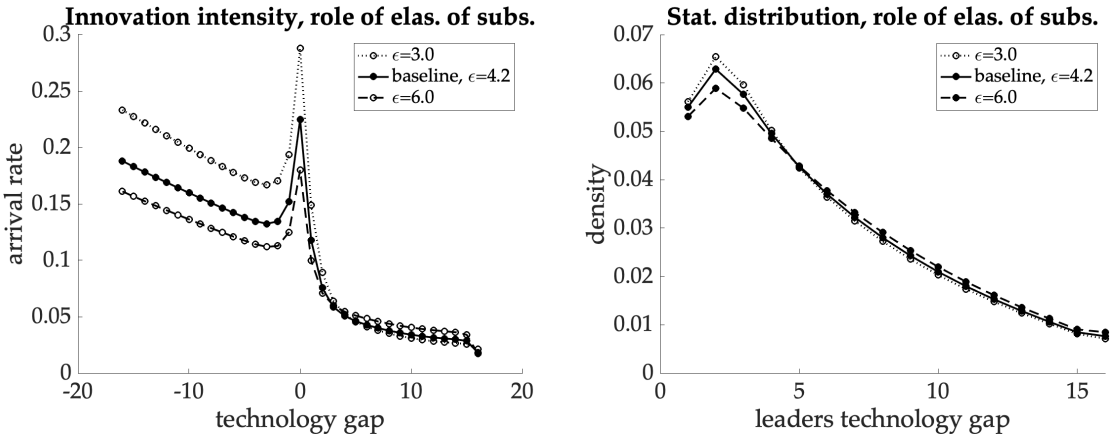

(a) Innovation policies **(b)** Stationary distribution

Figure B.4: Effect of the elasticity of substitution parameter ϵ on innovation policies x as a function of the technology gap for all firms (panel a) and the stationary distribution of sectors over the technology gap of the leader (panel b) under the calibration in Table 1. $\epsilon = 3$ is the market power experiment (dotted line), $\epsilon = 4.2$ is the baseline (solid line), and $\epsilon = 6$ is the superstar firm experiment (dashed line).

C Online Numerical Appendix

C.1 Solution Algorithm

For a given set of parameter values, the solution 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 using the fact that $\dot{v}_{mt} = 0$ on a balanced growth path. This process yields the optimal innovation policies of firms at each technology gap. Given the policy functions the stationary distribution of firms over technology gaps can be obtained by solving the system of equations described in section 3.5. To obtain the growth rate of output, solve the system described in appendix B.4. 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. To obtain micro-level moments, I simulate a discrete time version of the model with ten subperiods per year for a panel of 3000 firms for 400 years after the model reaches the steady state distribution over technology gaps.

C.2 Simulated Method of Moments Estimation

Let $M_j(\boldsymbol{\theta})$ denote the steady state value of moment j in the model as a function of the model parameters in vector $\boldsymbol{\theta}$. Let D_j denote the same moment in the data. The simulated method of moments estimation procedure seeks to find the vector of parameters $\boldsymbol{\theta}^*$ that solves:

$$\min_{\boldsymbol{\theta}} \sum_{j=1}^{J} w_j \left| \frac{M_j(\boldsymbol{\theta}) - D_j}{\frac{1}{2}M_j(\boldsymbol{\theta}) + \frac{1}{2}D_j} \right|$$

for J moments. I set the weights w_j such that the moments of interest, productivity growth and concentration, are weighted 5 times more than the other moments. I use the particle swarm optimization routine in Matlab to find the minimum.

C.3 Decomposition Table

Moment	Model		Data		Effect of each parameter						
	1990s	2000s	1990s	2000s	ρ	ϵ	η	δ_e	λ	α	ϕ
TFP growth	1.57	0.67	1.57	0.66	1.67	1.77	1.75	1.63	2.17	1.89	0.25
Leader market share	43.23	48.14	43.23	48.12	43.28	40.74	43.26	43.26	46.19	43.28	48.90
R&D/GDP	1.80	1.58	1.80	1.89	2.01	2.26	2.44	1.92	2.18	2.01	0.40
Profits/GDP	5.69	7.49	6.61	6.61	5.70	6.01	5.70	5.69	6.41	5.70	6.44
R&D/sales	4.63	4.14	3.75	4.06	5.19	5.50	6.31	4.97	6.03	5.17	1.00
Pat. qual., avg.	21.68	10.75	22.90	11.22	21.70	21.57	21.69	21.69	27.70	21.69	8.89
Pat. qual., followers	33.72	14.58	24.86	12.23	33.78	33.45	33.76	33.76	43.63	33.78	12.04
Pat. qual., leaders	9.99	6.94	10.00	5.17	9.97	10.04	9.98	9.98	12.28	9.97	5.76
Leadership turnover	13.93	9.44	13.74	9.28	14.12	14.50	14.35	13.21	14.32	14.67	9.93

Table 9: Effect of each estimated parameter change in Table 4 on the model steady state, holding other parameters fixed at estimated 1990s values. Decompositions follow the formula in equation 12.

C.4 Transition Dynamics

This appendix details the computational approach to solving the model's transition dynamics from one steady state to another. The approach is similar to Cavenaile, Roldan, and Schmitz (2021) but requires firm-level simulations at each iteration. I assume the economy begins in the initial (1990s) steady state in period $t = 1$ and arrives at the new steady state (2000s) by T, where T is large. The method is as follows:

1. Guess an interest rate path $\mathbf{r} = \{r_1, r_{1+dt}, r_{1+2*dt}, \ldots, r_T\}$.

2. Given the steady state values $v_{m,T}$ assumed at T, solve backward for innovation policies at $T - dt$ as:

$$x_{m,T-dt} = \begin{cases} \left(e^{-r_T dt} \frac{\sum_{n=m+1}^{\bar{m}} \mathbb{F}_{m,T-dt}(n)[(\lambda^{n-m})^{\frac{1}{\beta}-1} v_{n,T} - v_{m,T}]}{\alpha} \right)^{\frac{1}{\gamma-1}} & \text{for } m < \bar{m} \\ \left[e^{-r_T dt} \frac{1}{\alpha} (\lambda^{\frac{1}{\beta}-1} - 1) v_{\bar{m},T} \right]^{\frac{1}{\gamma-1}} & \text{for } m = \bar{m} \end{cases}.$$

3. Given the policy functions at $T - dt$ and the interest rate guess, solve for the value functions $v_{m,T-dt}$:

$$\begin{aligned} v_{m,T-dt} = & \left(\pi(m) - \alpha \frac{x_{m,T-dt}^{\gamma}}{\gamma} \right) dt \\ & + e^{-r_T dt} (x_{m,T-dt} dt \sum_{n=m+1}^{\bar{m}} \mathbb{F}_{m,T-dt}(n)[v_{n,T}(\lambda^{n-m})^{\frac{1}{\beta}-1} - v_{m,T}] \\ & + x_{-m,T-dt} dt \sum_{-m+1}^{\bar{m}} \mathbb{F}_{-m,T-dt}(n)[v_{-n,T} - v_{m,T}] \\ & + \delta_e dt(0 - v_{m,T}) + v_{m,T}). \end{aligned}$$

4. Repeat this backwards iteration for $x_{m,t}$ and $v_{m,t}$ until $t = 1$.

5. Initialize a panel of firms with stationary distribution μ_{mt} consistent with the stationary distribution in the initial steady state. Given $\{x_{m,t}\}_{t=1}^{T-dt}$, simulate the quality distribution forward to obtain aggregate R_t, K_t, Y_t and obtain household consumption from the resource constraint: $C_t = Y_t - K_t - R_t$ for $t \in \{t, \ldots, T\}$. Repeat this forward simulation S times.

6. Check if the conjectured interest rate sequence $\mathbf{r}$ is consistent with the resulting sequence of consumption growth rates using the households' Euler equation.

7. Update the guess of $\mathbf{r}$ to $\mathbf{r}_{new}$ using the average of the implied sequence of interest rates from the Euler equation over the S simulations.

8. Repeat until $|\mathbf{r}_{new} - \mathbf{r}| < \epsilon_{tol}$ for some small tolerance value ϵ_{tol}.

At each iteration I simulate a panel of 3,000 firms 30 times over 40 years with $dt = 0.02$.