Need for Speed: Quality of Innovations and the Allocation of Inventors

Santiago Caicedo * Jeremy Pearce†

March 2024

Work in Progress

Link to Most Recent Version

Abstract

This paper studies how the *speed-quality* tradeoff in innovation interacts with firm dynamics, concentration, and economic growth. Using microdata and a change in policy on patent duration, we document the existence of this tradeoff both in the aggregate and at the firm level. We show long-run trends in the increasing speed of innovation alongside declining quality at large firms, with the allocation of inventors playing an essential role. We develop a theoretical framework incorporating the speed-quality tradeoff and show that allocating less labor towards speed increases growth, particularly in the presence of private benefits to innovation. We estimate a quantitative endogenous growth model to study how firms' substitution across speed and quality interacts with aggregate outcomes. Quantitatively, the transition to faster, lower-quality innovations has a significant impact on growth, mainly through a shift in innovation production technology. We argue that even when the reallocation of inventors across speed and quality has modest effects on growth, endogenizing both the speed and quality decisions of firms is crucial for studying innovation.

JEL Codes: O30, 031, O33, J63

*Caicedo: Northeastern University, sa.caicedo@northeastern.edu.

†Pearce: Federal Reserve Bank of New York, jeremy.pearce@ny.frb.org.

We thank Isabella Becerra and Nicolás Fajardo for their excellent research assistance. The views expressed here are the authors' and do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

One striking pattern in modern innovation has been a trend in the increasing *speed* of patents produced while *quality* is declining.[1] As all firms increase the rate at which they patent, large firms, in particular, are experiencing reductions in their innovation quality. Given rising market concentration and the role of large firms as the major drivers of innovation outcomes (Klenow and Li, 2021; Aghion et al., 2023), this trend connects to essential questions on falling business dynamism, entry, and aggregate growth. Further, the tradeoff between speed (the "arrival rate") and quality (the "step size") is at the core of innovation. While the tension between speed and quality is recognized to play a major role in science (Hill and Stein, 2022), work on the interaction of this tradeoff with innovation, firm dynamics, and economic growth has been notably absent.

Insights on how firms pursue these strategies and how they impact competitors and the overall market are essential questions in an era of falling business dynamism. We unpack the implications of the speed-quality tradeoff by asking: what are the forces driving an increase in speed and reduction of the quality of patents of top firms, and how does this connect to falling business dynamism and the growth slowdown? We focus on this through the lens of how firms deploy inventors across projects, combining micro-data on inventors and firms disciplined by a novel theoretical framework. We start the discussion with three facts on the changing nature of innovation.

- **Fact 1: Speed.** Firms and inventors are producing patents at faster rates, with larger firms even faster than smaller ones.
- **Fact 2: Quality.** Large firms' patent quality is declining relative to small firms even as speed and concentration increase.
- **Fact 3: Speed $\times$ Quality.** From a change in policy in patent duration, we find that when firms increase speed in patenting, patent quality declines significantly.

Fact 1 illustrates the increasing speed of patent production, which is occurring at both the firm (patents per firm per year) and the inventor (patents per inventor per year) level. This is especially pronounced for large firms, which both produce patents faster and have a larger share of patents. However, large firms are not producing patents of higher quality. Fact 2 notes that large firms' patent quality is declining over time relative to small firms. This can be seen in the raw data and in event studies, where we observe inventors move from small to large firms and compare them to similar inventors in the prior period. Before 1995, they experienced an increase in quality. After

[1]Many have discussed declines in the marginal quality of innovation theoretically (Jones, 2023), and empirically (Kalyani, 2022). Another set of papers has noted trends in declining business dynamism (Decker et al., 2016; Akcigit and Ates, 2023 alongside many others). This paper introduces facts on speed and links it to the declines in quality.

1995, large firms induced a decline in quality relative to the previous period and the control group. The faster but lower quality patenting at large firms presents a tension on the role of speed and quality – is it a fixed factor of firms or endogenous? Fact 3 studies the endogenous speed-quality tradeoff through a policy change on the patenting duration we use to explore this tension. We study the introduction of a policy that induced firms to increase the speed of their patents to study how firms trade off speed against quality when it comes to the allocation of their inventors. We find that when firms are exposed to the "need for speed," they experience a decline in quality, indicating that the speed-quality tradeoff is a relevant endogenous choice within firms.

Guided by these empirical observations, we build a general theoretical framework centered around the speed-quality tradeoff. Our analysis highlights the importance of considering both speed and quality as endogenous factors, particularly in the presence of private benefits associated with innovation. These private benefits encompass any benefits independent of the innovation's quality, which can arise from various contexts, such as when firms profit from patent protection or have an external impact through creative destruction. Regardless of the origin or interpretation of private benefits, our main findings indicate that as these benefits increase, there is a shift in labor allocation towards faster innovations, leading to less innovation and growth than would otherwise be attainable.

Our framework's insights apply to many widely used innovation-led growth models. In the quantitative section, we build an endogenous growth model that bridges the micro-evidence on speed and quality with aggregate innovation and economic growth. Placing the speed-quality tradeoff at the forefront, we connect firms' microeconomic choices regarding innovation to macroeconomic outcomes, including concentration, overall innovation, and economic growth. To our knowledge, this is the first quantitative endogenous growth model with the tradeoff between speed and quality.

In our empirical analysis, we start by describing macroeconomic trends on the speed and quality of innovation across firms. We find that larger firms always produce faster patents, but as their speed has increased relative to small firms, their quality has declined. To further understand the dynamics of large and small firms, we perform event studies where inventors move from a small firm to a large firm. Comparing inventor movers to similar inventors who remain in small firms, we find that when inventors move to large firms, they produce faster patents. In the past, the quality of patents did not change significantly, though controlling for initial firm quality, we find that quality has declined in recent years. We further complement this evidence with more detail on market dynamics. We find that while relative patent quality measured by citations is decreasing for large firms, the private value is increasing. Further, in technology areas where large firms pursue higher-speed projects, there is less entry in the following period, whereas higher-quality patents encourage entry. Theoretically, this could lead to a disconnect between the socially optimal allocation and firm-level choice, a feature we discuss in our quantitative generalized model.

Before turning to the quantitative model, we explore the speed-quality tradeoff through a change in policy on the "need for speed". Is the speed and quality fixed at the firm level, or do firms make direct decisions that affect these outcomes? A 1995 policy change can help us answer this question and furnish our theoretical model with micro-evidence. We observe a natural experiment that induced exposed firms to increase the speed of their patenting over a 6-month period. In December 1994, it was announced that if firms successfully applied for a patent by June 7, 1995, their expected patent length was higher than if they applied by June 8, 1995. This effect was most pronounced for firms in specific fields that had long lags between application and grant date due to patent examiner time, as the length would be extended for firms with more than three years from initial application to grant. To measure the impact of the change, we classify firms based on their specialization across technological classes using the three-digit international patent classification (IPC3) system. The differential lag between patent applications and grant dates, as well as the varying private value of patenting across technological classes, affects the impact of the policy change for each firm. There is a noticeable hurry to patent that hits exposed firms and allows us to evaluate how their inducement to hurry shifted the portfolio of their patents. Given a fixed supply of inventors in the short run, this leads to lower-quality innovations and connects to the main tradeoff explored in this paper, which is the allocation of inventors across projects. We find that exposed firms' inventors produced 18% faster but 8% lower quality innovations. We further find, consistent with the discussion on the relation between speed and entry, persistently lower entry in the more exposed classes even a few years after the event.

Building on the empirical evidence of firms actively balancing speed and quality, we develop and estimate an endogenous growth model that embeds this tradeoff. We use the model to study the implications of adding the speed and quality decision, showing their relevance by replicating the observed rise in patenting concentration between the 1980s and the 2010s, which has been emphasized in recent literature (Akcigit and Ates 2021, 2023). We estimate the model to capture key moments: the growth decline over two periods, faster patenting and innovation concentration, lower patent quality among large firms, declining innovation by new entrants, and an increase in the private value of patents relative to their quality.

Leveraging our estimated model, we perform two main quantitative exercises. First, we decompose the contributions of speed and quality to the decline in growth and document the changes in the allocation of inventors. Second, we quantify the impact of labor reallocation on innovation and growth, evaluating the relevance of endogenous quality. The decomposition reveals the increasing importance of incumbents in driving growth. Faster patenting by incumbents partially compensates for the fall in the overall quality of innovations. Inventors' allocation shifts towards speed, rising from 74% to 86% of total labor, with a substantial decrease in the proportion dedicated to quality. We show the shift in labor reallocation is a significant factor contributing to the decline in patenting quality. In our second quantitative exercise, we explore the effect of reallocating labor on innovation and growth. Despite the significant increase in labor allocated to speed, changes in innovation production technology attenuate the effects on growth. Reallocating labor across speed and quality to maximize growth yields modest improvements, raising growth from 1.32% to 1.34%. This is not to say that endogenizing quality does not matter. To illustrate this point, we compare an economy with the same estimated parameters but with fixed quality. We show that as the private benefit increases, the predicted trend and level of growth diverge substantially.

Our ongoing analysis aims to extend the model to explicitly consider the effect of the private benefit on firm entry and, more generally, characterize the conditions for which the speed and quality tradeoff matter most.

Related Literature. This paper speaks to some of the fundamentals of innovation across firms, with firm heterogeneity, inventor allocation, and market concentration. In doing so, our paper combines three fields: i) firm dynamics, innovation, and economic growth; ii) the importance of human capital in innovation and economic growth; and iii) a primarily empirical literature on the mechanics of innovation and science of science.

We start by addressing our contribution to the firm dynamics and endogenous growth literature. There is an extensive body of research that has emphasized the role of heterogeneous innovations in growth and firm dynamics (Grossman and Helpman (1991);Aghion and Howitt 1992; Klette and Kortum 2004; Lentz and Mortensen 2008; Akcigit and Kerr 2018; Garcia-Macia et al. 2019; Peters 2020). Some recent papers in this literature have used a firm dynamics framework to study the rise in market concentration and productivity slowdown (such as Akcigit and Ates, 2021, 2023 and Aghion et al., 2023). These papers make the point that market concentration will have important interactions with innovation, and these two concepts cannot be considered separately from each other. The role of heterogeneous innovation strategies, a central component of this paper, addresses how concentration feeds back to choices over innovation, a discussion that connects to classic questions over competition and innovation (Aghion et al., 2005).

Related to this vast, mostly theoretical literature is a growing set of studies focusing on falling business dynamism and rising concentration (Decker et al., 2016; De Loecker et al., 2020; Covarrubias et al., 2020). Liu et al. (2022) link rising concentration and falling growth to interest rates. There are various arguments about the core sources driving this rise. Eggertsson et al. (2018) discuss the rise of monopoly power and Autor et al. (2020) discuss the rise of superstar firms. Bessen (2017), Crouzet and Eberly (2018) and De Ridder (2019) focus on technology and intangibles. Our framework focuses on the innovation side and evidence of declining innovation impact and growth, as in Bloom et al. (2020), Akcigit and Ates (2021), Kalyani (2022), and Aghion et al.

(2023). To understand this, we direct particular attention to human capital and the endogenous speed-quality tradeoff at the firm level. We demonstrate that the distribution of innovation across firms and the tradeoffs they face shift the overall composition of innovation.

For firms, inventors are the core ingredient of their innovation. This connects to a growing literature within economic growth that puts human capital at the center. We stress two reasons for this decision. First, human capital is a central ingredient to economic growth. Many economists have noted this point (Mincer, 1984; Lucas, 1988). More recent work has shown human capital to be a central force for growth empirically, theoretically, and quantitatively (Waldinger, 2016; Akcigit et al., 2017; Jaimovich and Rebelo, 2017; Akcigit et al., 2018, 2020). Second, we observe this central input in innovation in the data, where we observe individual inventors, firms, and their innovative output. By tracking individual inventors in different firms, we can speak to different aspects of the innovation process.

In presenting human capital and innovation at the center, we stress how inventors are the primary input to innovation, yet firms determine the direction and deployment of innovations. Firms determine the environment inventors work in, the projects they work on, and the managerial feedback they receive. Firm management facilitates human capital interaction, leading to learning and growth (Garicano and Rossi-Hansberg, 2006; Lucas and Moll, 2014). The firm structure can even be thought of as a technology for producing innovation and production output (Bloom et al., 2016). At the center of the firm's choice is the allocation of talent to project type, which also speaks to an important literature on the sorting of individuals to occupations and tasks (Willis and Rosen 1979; Murphy et al. 1991; Hsieh et al. 2019). The task-skill match is essential for individuals to realize their potential and the link between talent and output. The direction and type of innovation firms allocate their inventors to will shape this force significantly.

When linking inventor allocation to firms, our paper speaks to recent literature on talent allocation across firms and economic growth (Manera 2022; Ma 2022). Our study is an early contribution to the interaction between inventor allocation and innovation type. As such, we build on a literature that has noted the interaction between firm size and innovation types (e.g., Acemoglu and Cao, 2015 and Akcigit and Kerr, 2018). We bring this literature closer to the discussion on human capital, which is essential to understanding the links between inventor allocation, firm dynamics, and economic growth. This paper connects to a rich literature on both dimensions yet fills a gap in their interactions.

Lastly, in studying the mechanics of innovation, we build mechanisms related to the micro-foundations of innovation. Firm choice over inventor composition will naturally shape the type of projects firms pursue and the overall nature of innovation. For instance, Wu et al. (2019) points out that large teams are much more likely to develop existing fields, whereas small teams pursue more radical ideas. In this paper, we note that large firms are more likely to deploy large teams.

Large teams may run into issues regarding splitting the rents (Kline et al., 2019), but also build team-specific capital, which matters for impact over time (Jaravel et al., 2018). This is especially pronounced for the interaction with superstars (Azoulay et al., 2010). Yet, it is also important to note teams sort endogenously depending on the returns to innovation and individual expertise (Pearce, 2020). This endogenous mechanism is driven by a firm's returns in our paper, and thus, the composition of firms will shape the composition of teams and innovation.

One central point of our paper is that firms may prioritize speed or quality depending on the tradeoffs in innovation. This connects to a broad question on allocative efficiency in innovation. Yet, in noting this tradeoff, our mechanism builds classic quality-quantity debates that have long interested economists (Becker, 1960; Becker and Lewis, 1973) to the domain of innovation and growth. Indeed, Hill and Stein (2022) find important tradeoffs come into play when teams are incentivized to produce more quickly and focus on speed, or quantity, over quality. Within this field, we highlight how different innovation strategies may have different spillovers depending on their effects on existing markets. This connects to the seminal creative destruction frameworks of (Aghion and Howitt, 1992).

The structure of the rest of the paper is outlined as follows: Section 2 outlines our theoretical framework on the speed-quality tradeoff. Section 3 presents empirical evidence on the changes in the speed and quality of innovations and inventor allocation and the policy change in patent duration. In Section 4, we describe the estimation of a quantitative version of the model, and we outline counterfactual exercises to study the relationship between innovation, inventor allocation, and growth.

2 Framework: Speed and Quality Tradeoff

We start by examining a theoretical framework highlighting the tradeoff between speed and quality in innovation. To illustrate the primary mechanism, we focus on firms' innovation problem, which is at the core of most innovation-led endogenous growth models. We extend the basic framework in these models by considering innovating firms balancing between allocating labor to increase the speed of innovations (arrival rate) or their quality (step-size). We are interested in understanding how innovations' value and profitability affect labor allocation across the speed and quality margins. In Section 4, we embed the speed-quality framework into a fully-fledged endogenous growth model that extends the work of Acemoglu and Cao (2015). This particular choice of the model for our quantitative exercises is useful for matching some key trends in the data, especially when targeting the high patenting concentration. However, the speed-quality tradeoff applies more broadly to other workhorse growth models, from Aghion and Howitt (1992) to Akcigit and Kerr (2018), with qualitatively equivalent results.

Suppose firms choose to allocate labor between speeding up the arrival rate of innovations, $x(l_x)$, and improving the quality of their innovations, $Q(l_q)$. The value of a successful innovation $V(q)$ depends on the technological level q of the product line firms innovate on, although this is not essential for the results. Firms obtain a flow profit $\pi(q)$ from getting an innovation in product line q and producing the corresponding goods in the economy. Suppose the flow profit is non-decreasing and weakly concave in q. In addition to this flow profit, firms can also have a private benefit of innovating, $B \geq 0$. Conceptually, this private benefit captures any benefit of increasing patenting speed that is not directly related to the quality of the innovations. This benefit could have endogenous and exogenous components depending on the specific model. For instance, B can be viewed as reflecting the protective value of patents or the value of creating patent thickets around an initial innovation. On the other hand, in models with external innovations, B also endogenously captures the 'business stealing effect' and franchise value from getting a product from another firm. Here, we are agnostic about the particular nature of B, but instead examine the consequences of having this private benefit for innovating firms' speed and quality decisions. We show that having the private benefit is fundamental for the positive and normative implications of the speed-quality tradeoff, as it changes the allocation of resources for innovation and affects innovations' private and social value.

The firms' innovation decisions in continuous time can be generically described by the Hamilton-Jacobi-Bellman (HJB) equation,

$$r_t V_t(q) - \dot{V}_t(q) = \pi_t(q) + \max_{l_x, l_q} \left\{ x(l_x) \left[V_t(q + Q(l_q)) - V(q) + B \right] - w_t(l_x + l_q) \right\} - \tau V_t(q), \tag{1}$$

where τ denotes the rate at which other firms innovate on the product line with quality q (creative destruction), r_t the interest rate, and w_t the wage rate.

Suppose the arrival rate and quality functions are increasing, jointly concave, labor is essential and V is differentiable.[2] This guarantees an interior solution where the optimal labor demands imply the marginal value of speed and quality equates to the marginal cost of labor,

$$\begin{aligned} x'(l_x) \left(V_t(q + Q(l_q)) - V_t(q) + B \right) &= w_t, \\ x(l_x) V'(q + Q(l_q)) Q'(l_q) &= w_t. \end{aligned} \tag{2}$$

Based on these first-order conditions (FOCs), we show two results: first, having the additional private benefit $B > 0$ leads to increases in the labor allocated to speed, and second, the increase in speed implies an allocation that does not maximize growth.

[2] Mathematically, suppose the arrival rate satisfies, $x'(\cdot) > 0$, $x''(\cdot) < 0$, $x(0) = 0$; the quality of innovations $Q'(\cdot)$, $Q''(\cdot)$, $Q(0) = 0$ and the determinant of the Hessian associated to the expected value of innovation is positive, $x''Q''(AQ + B)x > (x'Q')^2$. In the quantitative model we consider in Section 4, $V(q)$ is differentiable.

Formally, let l_x^*, l_q^* denote the labor demands for speed and quality that solve (2). We can show that the relative demand for speed is increasing in B. All proofs are in Appendix B.

Proposition 1 (Labor Allocation)**.** *The labor for speed relative to quality, $\frac{l_x^*}{l_q^*}$, is increasing in B.*

Proposition 1 offers a simple framework to study the variations in speed and quality during the firm's life cycle and over different periods. Changes in B may indicate changes in the market conditions in which firms operate or the nature of the spillovers in innovation. For example, protective patents may increase B, augmenting the profitability of innovations, irrespective of their quality. On the other hand, changes in the private value of innovation can also stem from firms not fully internalizing the spillover effects of innovation. They can lead to a socially suboptimal level of innovation.

Naturally, the allocation of inventors across the speed and quality margins can also affect growth. The expected contribution of innovating firms to growth is given by the product of the arrival rate times the quality of the innovation: $x \times Q$. The contribution is maximized when the marginal use of labor is equalized across the speed and quality margin,

$$x'(l_x)Q(l_q) = x(l_x)Q'(l_q).$$

Comparing this condition to (2) implies that whenever there is a private benefit, the growth contributed by innovation firms would be larger if less relative labor is allocated to speed.

Proposition 2 (Growth)**.** *If $B > 0$, less relative labor to speed increases growth from innovating firms, xQ.*

Together, these elementary propositions characterize the effect of endogenizing quality on labor allocation and growth. An additional private benefit implies incumbent firms will over-invest in speed at the expense of quality. This leads to the misallocation of innovation resources, in this case inventors. In Section 4, we develop a quantitative version of the model and take it to the data to estimate the private benefit B needed to rationalize the observed trends in the speed and quality of innovations across firms of different sizes and over time. Before detailing the full quantitative model, we empirically document the speed and quality tradeoff at the aggregate and firm levels.

3 Data and Empirical Analysis

This section presents the USPTO patent data and details the empirical evidence that serves as the motivating evidence for our study. The results in this section suggest that firm and innovation heterogeneity play a central role in our analysis. Large and small firms have heterogeneous rates and types of innovation. We also show evidence that the speed-quality tradeoff also happens within firms and types of innovation. Section 3.1 discusses the data development; Section 3.2 presents the macroeconomic trends on speed and quality across firms; Section 3.3 focuses on the change in the speed and quality of projects as inventors move across firms; Section 3.4, we study a policy change in patent duration to provide quasi-experimental evidence showcasing the speed-quality tradeoff within innovations at the firm level.

3.1 Data and Definitions

Patents have emerged as a cornerstone of observable objects that enable economists to study trends in innovation. When individuals come up with a new and useful idea, they can protect their property rights on the exclusive use of this idea through the patent system. We aim to understand innovations through this lens.

We use USPTO patent data compiled by PatentsView (USPTO 2019), which provides disambiguated information on inventors and firms. This data allows us to track the careers of inventors and firms over time, utilizing patents granted between 1975 and 2015. The dataset provides, for each patent, an identifier of the primary assignee (e.g., firm) who owns the patent, alongside the "true and only inventors" who were the individuals behind the patent. Alongside this information, there is rich detail on location, technology class, citation flows, and patent application date and grant date.

To compute the social value patent, we use citation data from the USPTO. To compute the private value of patents, we use the stock value data from Kogan et al. (2017). We categorize firms as large or small based on the number of patents produced in a year, with the top 1% firms being defined as large and the bottom 99 % being defined as small. The concentration of innovation is measured using this categorization. Patent quality is measured using the logarithm of forward citations controlling for IPC3 technology class and year. This measure is widely accepted in the patenting literature as a meaningful indicator of the social value of patents (Hall et al. 2001; Akcigit and Kerr 2018). To compute the speed of patents, we take the average time at the *inventor*-level between patents, to proxy for the time it takes to produce the focal patent. Internal innovations are defined as patents within the same technological class as a previous patent by the same firm, while external innovations are defined as the firm's first patent in a new technology class.

One aspect of the data that is essential to our study is the ability to track inventors over time. PatentsView's disambiguation enables researchers to view inventors as they move across firms and locations and speak to changes in citations. Most novel to this project is introducing data on the time between patents at the inventor level. We simply take the "time-to-build" for an individual-patent pair as the number of years between an individual's patents. Fast patents have less time between them at the *inventor* level. This proxy for the fact that we consider inventors to be the main scarce input in idea production and track their time between patents as roughly indicating how long it takes to produce an idea.

We present the summary statistics in Table 1. We split the data into aggregate variables, firm-level variables, and patent-level variables. At the aggregate, we report the total number of patents, inventors, and firms; at the firm level we report patents at the firm level, patents per year, and share of the largest firms; at the patent level, we report averages of citations, private value, time between patents, and team size.

Table 1: Summary Statistics

— *Panel A. Aggregate* —			
Variable	**Value**		
Sample Period (Granted)	1976-2015		
Unique patents	5,298,348		
Unique inventors	3,004,814		
Unique firms	368,287		
— *Panel B. Firm Level* —			
Variable		**Average**	**SD**
Patents per Firm	-	14.4	383
Patents per Year	-	1.70	11.4
Top 1% Firm Share of Total Patents	-	42%	-
Top 10% Firm Share of Total Patents	-	66%	-
— *Panel C. Patent Level* —			
Variable	**Obs**	**Average**	**SD**
Forward Citations	5,298,348	1.97	5.89
Patent Value ($ Millions)	1,281,667	21.96	70
Inventor Time-to-Patent (Years)	5,040,543	1.54	2.28
Team Size	5,298,348	2.59	1.81

Notes: The full sample from PatentsView runs from 1976-2015. We take 1980-1985 and 2010-2015 as our benchmark sample periods.

We focus on a couple of simple takeaways from Table 1. First, there are many patents (more than 5M), of which around 25% (1.3M) are linked to publicly traded firms with public valuations. The majority of patents have a time-to-patent variable (which is missing if an inventor only has one patent). Finally, the overall shares of the top 1 and 10% are large, but growing over time. Lastly, some variables (e.g., citations, private value, and time to patent) exhibit large variation, as the standard deviation is significantly larger than the mean. The overall takeaway is that this data provides significant variation in firm size distribution, patent quality, and time to produce patents.

We will use these variables to study the role of the choice between speed and quality in addition to firm heterogeneity in the rest of this section.

3.2 Macro Trends: Concentration and Quality

This section studies the trends in patenting speed and quality and the role of firms. We find that there has been a strong trend in the increase in patent speed for both large and small firms, while large firms' patent quality is decreasing relative to small firms. The long-run shifts in speed and quality, alongside its interaction with human capital, provide important facts that relate to our theoretical framework from the previous section and discipline our quantitative exercises.

Figure 1 presents evidence of speed and quality jointly. Figure 1a shows the steady increase in the number of patents per firm since the late 1980s. The growth in patenting has been more pronounced for top 1% firms, leading to more concentration in innovation. While large firms produce faster patents, this has been accompanied by a decline in the quality of patents produced by these firms. We split firms into the bottom 99% and top 1% of US and foreign corporate firms in terms of number of patents. Large firms now produce innovations with less social value, as shown by our measure of forward citations. Figure 1b illustrates this decline, which holds true even when using other measures of patent quality, such as different lags for citations, using raw variables and different sets of controls, excluding forward self-citations, considering "breakthrough" patents (top 5% and top 1% most cited) or looking at the generality and originality of patents (see Figure A.1 in the Appendix).

Figure 1: Quantity and Quality at Bottom 99% and Top 1% Firms

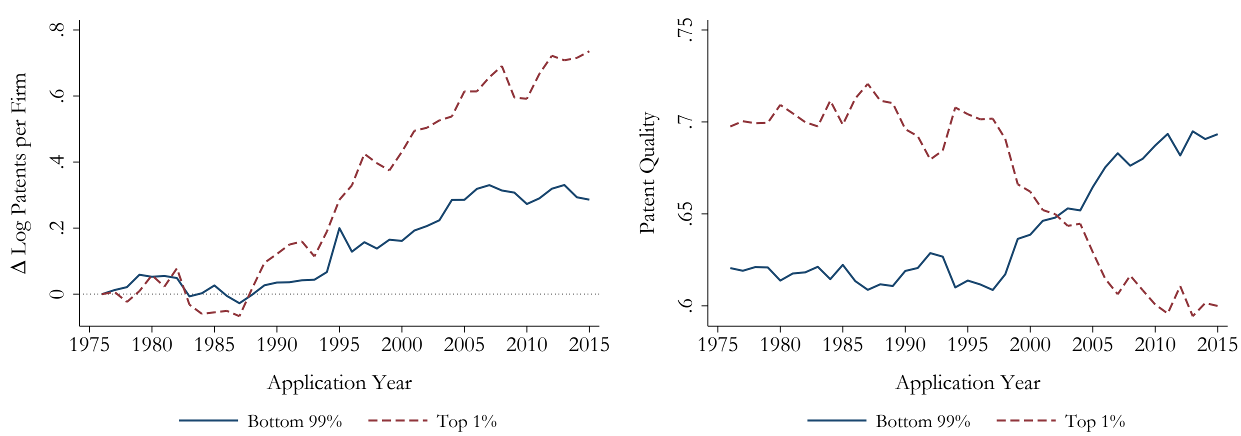

(a) Quantity: Log Patents per Firm

(b) Quality: Log Citations

Notes: Quantity: log number of patents per firm, normalized to equal 0 in 1976. Patent quality=$Log(1 + 3yr\ forward\ citations)$, residualized using time, technology class, and team size controls. The sample is split into the top 1% and bottom 99% in cumulative patent stock.

One salient trend in the patent data is the increase in patenting speed at the inventor level. Prior to 1995, firms produced patents at a fairly constant clip, with the average inventor producing a patent about once per 1.35 years. This number steadily declines after the mid-90s, and in the last five years of the sample, 2010–2015, inventors produced slightly more than a patent per year. This increase in patent rates occurred for both small and large firms (Figure 2).

Figure 2: Speed– Inverse Time-to-Patent at Inventor Level

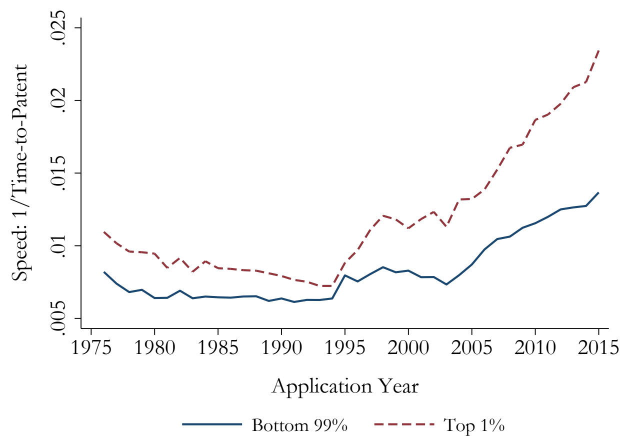

Notes: Time-to-Patent: number of days since the last patent, measured at the inventor level.

The above trends indicate a speed-quality tradeoff over time at the macro level, especially for large firms. Large firms produce faster patents and of relatively lower-quality in the most recent decades. This becomes potentially concerning as larger firms are taking larger market share, which may drive down overall innovation. To dig deeper into the microeconomic evidence of this fact, we turn to studying inventors moving across firms.

3.3 Movers: Speed and Quality at the Inventor-Level

The macro trends are important to understand firm dynamics, but recent work has repeatedly pointed out that human capital is the central input to the innovation process (Waldinger, 2016; Akcigit et al., 2018, 2020). Further, individuals provide a data-rich way to understand firm choice in innovation, as individuals who work at different firms can be observed over time. We start by showing that the increasing quantity of innovations at large firms is in part associated with an increase in *speed* at the *inventor*-level.

To shed light on the interaction between human capital and firm innovation, we evaluate events when inventors move across firms. In particular, we compare inventors moving from small to large firms with those moving from large to small firms. We match pairs and control for "mover" specific measures to provide more insights on the links between the microeconomics of innovation and how innovation production has changed over time.

Event Study: Move from Small Firm to Large Firm— To complement our previous analysis, we study how inventor patenting changes when moving from small to large firms before and after 1995. We chose 1995 as the midpoint of our sample due to some of the changes in macro trends that started in the mid-1990s (see Figures 1 and 2). We find that, in contrast to the period before 1995, inventors produced faster but lower-quality patents after 1995 when moving to large firms.

For this analysis, we match inventors working at small firms with similar citation patterns and study the changes in speed and quality for an inventor moving from the bottom 99% to the top 1% of firms, with a control group of inventors who remain at small firms. More formally, the regression model takes the following form:

$$y_{i,j,t} = \alpha + \sum_{a=-3}^{5} \beta_a \times \mathbb{I}\{t = a\} + \sum_{a=-3}^{5} \zeta_a \times move \times \mathbb{I}\{t = a\} + \Lambda_i + \Gamma_t + \Sigma_j + \epsilon_{i,j,t}, \tag{3}$$

where $y_{i,j,t}$ denotes the dependent variable of individual i at firm j at time t, where time t indicates the patent number around the time of the move. We study four outcomes: the logarithm of the number of patents at the firm level, the logarithm of speed measured as the inverse of the inventor time between patents[3], patent quality measured by the residualized logarithm of 3-year forward citations controlling for time, technology class, and team size fixed effects, and finally, the measure of the patent value from Kogan et al. (2017). Λ_i is an individual fixed effect, Γ_t is an age fixed effect, and Σ_j is a fixed effect controlling for the firm's initial quality (as in initial patent citations). These fixed effects help simply identify the marginal effect of moving from a small firm to a similar, but larger, firm. The coefficients ζ_a isolate this effect.

Figure 3 highlights significant differences in the speed-quality tradeoff between the two periods. In both periods, we observe that inventors who move to large firms tend to produce faster patents (at both the firm-level arrival rate, Figure 3a, and the inventor speed, Figure 3b). Furthermore, the firm-level arrival rate also increases in the later period, indicating the relative increase in patenting for larger firms, even with the same inventors. However, after 1995, we also find that moving to a large firm led to a decrease in the quality of the produced patents (as shown in Figure 3c). The quality-quantity tradeoff seems to become more pronounced for large firms.

Notably, Figure 3d shows the value of new innovations produced in large firms is larger than in small firms, and this trend has remained consistent over time. This suggests that the decline in patenting quality does not necessarily result in a corresponding decrease in patent value for large firms. It is profitable for large firms to produce faster, lower-quality patents.

Our empirical analysis suggests that the allocation and use of inventors have changed, with large firms using inventors to produce lower-quality inventions. These trends raise the possibility of misallocation of inventors. However, these observations are only correlations. In the following

[3] We take the average time between patents of inventors on a focal patent to measure the speed of patenting.

Figure 3: Inventor Moves From Small to Large Pre-1995 vs. Post-1995

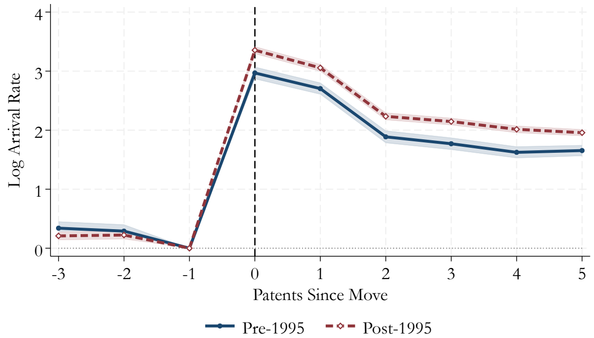

(a) Arrival Rate: Firm-Level

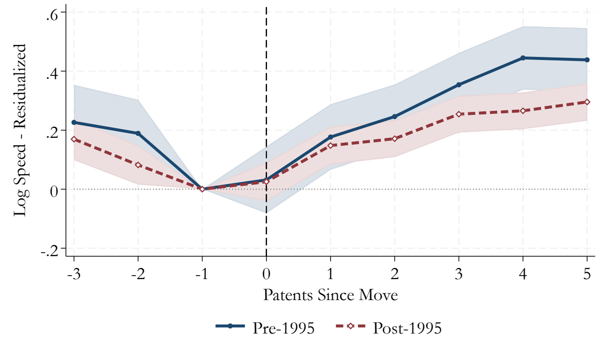

(b) Speed: At Inventor Level

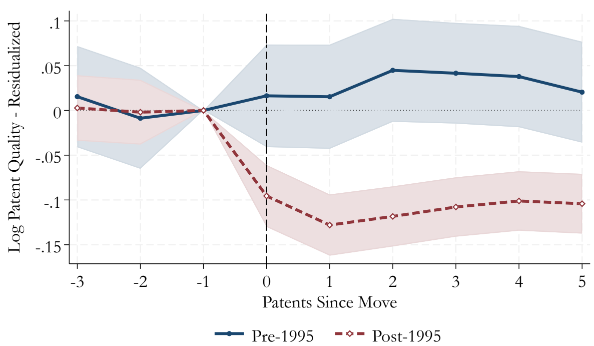

(c) Quality: Log Citations

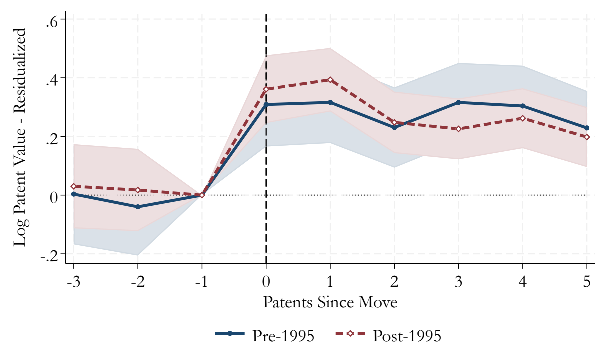

(d) Private Value: Log Stock Market Value

Notes: We measure the arrival rate at the firm level as the log number of patents per firm, the speed at the inventor level is the inverse of the average time to patent, patent quality=$Log(1 + 3yr\ forward\ citations)$, residualized using time, technology class, and team size controls and private value is the log of the stock market value change from Kogan et al. (2017). The regression model is specified in equation (3).

section, we provide further evidence from a policy change that sheds light on how firms adjust the speed and quality of their patents. This variation is also helpful for identifying how firms allocate labor to produce faster or higher-quality innovations, which is crucial for our quantitative analysis.

3.4 A Natural Experiment on the Speed-Quality Tradeoff in Innovation

The long-run trends in innovation suggest an important tradeoff between declining quality and increased speed of innovation, with inventor allocation playing a key role. In looking at this process, it is natural to ask whether this is a firm-level choice variable or a function of broader economic trends or the firm life cycle. To better isolate how much endogenous choices determine this tradeoff, we leverage a natural experiment in 1995 that exogenously sped up the production of innovation. We find that the allocation of inventors to speed and quality is a relevant margin for the firm and will later use this margin to study the elasticities of innovation speed and quality with respect to the number of inventors at the firm.

There are natural challenges to identifying the relationship between the quality of patents produced and the time it takes to patent them. Recent evidence on the time tradeoff in science points to the fact that faster projects tend to produce ideas of lower quality (Hill and Stein, 2022). However, this context does not make use of exogenous variation and may not extend to the inventor context within firms. Given our earlier evidence that large and small firms make different use of inventors of heterogeneous talent, this information is crucial to informing our policy counterfactuals. We make use of a policy change in order to identify this tradeoff at the firm and patent levels.

In our experiment, we study the impact of an exogenous shock to patent application deadlines on patent quality and speed to understand this tradeoff. This will inform our policy counterfactuals, as a change in the composition of firms will lead to changes in the allocation between high-speed and high-quality patents. This experiment has been studied to explore the nature of policy changes and how they interact with firm innovation levels (Abrams, 2008; Bertolotti, 2022). Our paper is the first we are aware of that uses this event to understand the dynamics of the faster production of innovation in response to the policy announcement.

On December 8, 1994, Bill Clinton signed the "Uruguay Round Agreements Act," which changed the structure of patent law. The change in law was to go into effect on June 8, 1995, and patents were to last 20 years from their application date rather than 17 years from the grant date. In the intervening period, firms that submitted patents would get the corresponding maximum amount of time. For example, if a patent takes 3.5 years from the application date to the grant date, the patent term would last 17 years from the grant date. If a patent takes 1.5 years from the application date to the grant date, the patent term will last 20 years from the application date. Thus, firms get an advantage from submitting a patent on June 7, 1995, instead of June 8, 1995, simply for the option value of the maximum time length. We use the fact that different firms have different exposures to this shock to understand the differences in the responsiveness of shocked firms to the policy change.

We study the differential effect of policy on the production of patents at firms that are most exposed (e.g., in the top quartile technology class in terms of lag from application date to grant date) and least exposed (e.g., in the bottom quartile technology class in terms of lag from application date to grant date). The change in policy leads to changes in firm patent production around the time of the event. Figure 4 plots the responsiveness of patenting overall around the time of the event and the different responses depending on technologies that are more exposed versus less exposed.

This change in patent applications around the time of the event is striking. More exposed firms had around three times as many patents in the second quarter of 1995. Outside of this period,

Figure 4: Patent Applications and Exposure

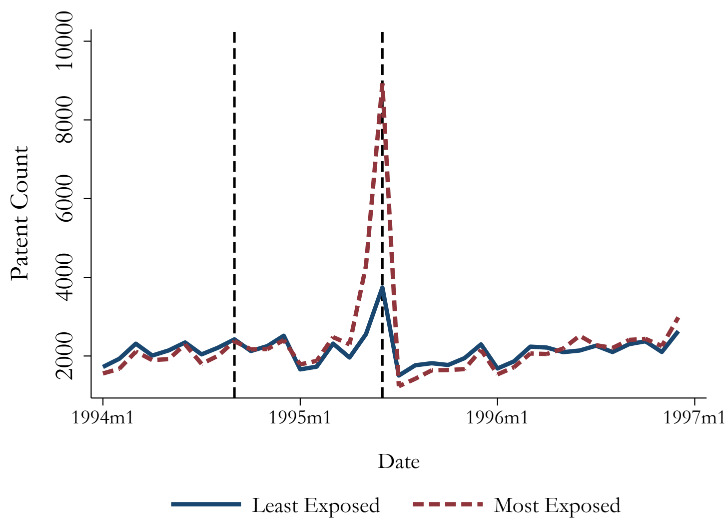

Notes: Least exposed: below 25th percentile and Most exposed: above 75th percentile by firm-technology class exposure.

the unexposed and exposed firms produce the same number of patents. With this shock, we can understand how firms pursue different strategies in the speed-quality tradeoff. By tracing out the responsiveness of unexposed and exposed firms, we can learn about the role of the speed-quality tradeoff– do firms sacrifice quality to increase speed? To do so, we evaluate at the patent level the changes in speed and quality between the event and announcement date. Since the average patent takes about 1.5 years to produce (as measured by the time between inventors working on a given patent), the 6-month window following the announcement date may cause some projects to be expedited, which may impact the underlying project quality.

We add controls for technology, firm, and date to control the relevant variables to observe changes within firms and technologies. This now captures within-firm variation depending on the exposure to policy change. We present the general framework of an outcome $y_{i(k)jt}$ for patent i in technology k and firm j at time t in Equation (4):

$$y_{i(k)jt} = \beta_0 + \beta_1 \Lambda_t \times exposed_j + \Lambda_t + \Gamma_j + \varphi_k + \epsilon_{ijt}, \tag{4}$$

where the outcome of interest is the interaction between the time variable and the firm exposure $\Lambda_t \times exposed_j$. This includes general controls for time (Λ_t), firm (Γ_j), and technology class (φ_k).

Figure 5: Speed and Quality

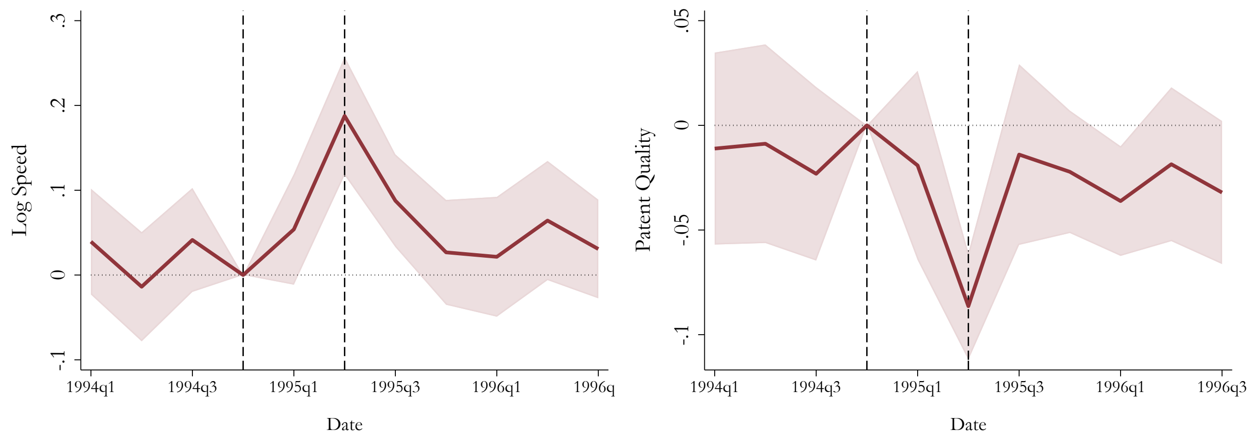

(a) Speed: Log(1/Time to Patent)

(b) Patent Quality Q

Notes: Coefficient plot of exposure indicator (above p75 vs. below p25) interacted with quarterly dummies. Patent level regression with firm and quarter fixed effects. Errors clustered at the technological class level (IPC3).

Figure 5 shows more exposed firms experience both a decline in quality and in the time needed to produce a patent. To produce faster patents, exposed firms experienced a reduction in quality. On average, exposed firms see an 18% increase in speed which is coupled with an 8% decrease in quality. This is starker for firms that did not increase their inventive labor (e.g., kept human capital stock fixed), as, on average, more exposed firms poached inventors to increase speed. These results point to two important and central messages of this paper. First, there is a speed-quality tradeoff in innovation, and depending on the spillovers, this can affect the market structure. Second, the use of human capital is a central ingredient for innovation and the key factor to understand this tradeoff.

In Appendix A, we further explore the mechanics of the responsiveness to this policy. We find that most of the responsiveness is due to "internal" patents of the firm–e.g., in their focal technology classes. Furthermore, this response is uniform across small and large firms, indicating both that the speed-quality tradeoff is relevant across the firm size distribution and that the shock can be understood in a somewhat uniform way across firms.

4 Quantitative Model and Application

Having established empirical evidence on firms actively balancing the speed and quality of innovations, in this section, we develop and estimate a quantitative endogenous growth model that embeds the speed-quality tradeoff from Section 2. More precisely, we extend the work of Acemoglu and Cao (2015), adding the speed and quality decision for incumbent firms and use the estimated model to match the trends in the increase in patenting concentration between the 1980s and the 2010s, which has been highlighted in recent literature (Akcigit and Ates 2021, 2023). We consider two main quantitative exercises: i) breaking down the contributions of speed and quality to the decline in growth and ii) quantifying the impact of labor reallocation on innovation and growth. To motivate these exercises, we revisit the evidence of the increase in the concentration of innovation and the allocation of inventors.

4.1 Growth Model With Endogenous Speed and Quality

Suppose an economy with a mass L of workers that provide raw labor for the final good production and a mass L_I specialized workers that produce innovations in the research sector. A perfectly competitive sector produces the final good, using the raw labor and a fixed continuum of intermediate goods k_j with $j \in [0,1]$, of quality q_j,

$$Y(t) = \frac{1}{1-\beta} L^\beta \int_0^1 q_j(t)^\beta k_j(t)^{1-\beta} dj.$$

Let $p_j(t)$ denote the price of the intermediate good j in period t. Solving the final good producers' problem implies an isoelastic inverse demand function, $p_j(t) = L^\beta q_j(t)^\beta k_j(t)^{-\beta}$. Below, we omit the explicit time t notation to reduce the clutter whenever there is no confusion.

The intermediate good is produced by a continuum of firms using linear technology. To make k_j units of intermediate good j, they pay a constant marginal cost ψ denominated in terms of the final good. To simplify the analysis, we assume that intermediate good producers with the highest quality are the monopoly producers of a variety.[4] Every period intermediate firms solve,

$$\Pi_j = \max_{k_j} \{p_j k_j - \psi k_j\} = \max_{k_j} \left\{L^\beta q_j^\beta k_j^{1-\beta} - \psi k_j\right\}$$

The solution implies $k_j = \left(\frac{1-\beta}{\psi}\right)^{1/\beta} Lq_j$, so profits are $\Pi_j = \beta \left(\frac{(1-\beta)}{\psi}\right)^{\frac{1-\beta}{\beta}} Lq_j =: \pi_I q_j$.

Let $\bar{q}(t) := \int q_j(t) dj$ denote the aggregate productivity of the economy integrating over the active product lines. Hence, the aggregate output in the final goods sector is,

$$Y(t) = \frac{1}{1-\beta} \left(\frac{1-\beta}{\psi}\right)^{\frac{1-\beta}{\beta}} L\bar{q}(t) \tag{5}$$

Innovation In our model, incumbent firms choose to allocate labor between speeding up innovation production (l_x) and improving the quality of their innovations (l_q). Skilled labor from

[4] See Acemoglu and Akcigit (2012) for a microfoundation. Similar results follow if we impose limit pricing.

inventors is required for both types of innovations. We allow firms to have an additional private benefit of innovating, $B \geq 0$, that is not affected by the quality of the innovations.

Incumbent firms innovate on their product line, q_j. They are exposed to creative destruction from entrant firms, which happens at a rate τ. The value function of an incumbent firm satisfies the HJB equation,

$$rV(q_j) - \dot{V}(q_j) = \pi_I q_j + \max_{l_x, l_q} \{x(l_x)[V(q_j + Q(l_q)) - V(q_j) + B] - w(l_q + l_x)\} - \tau V(q_j) \tag{6}$$

The optimal labor demands of an interior solution imply the marginal value of labor for speed and quality equates to the marginal cost of labor,

$$\begin{aligned} x'(l_x)(V(q + Q(l_q)) - V(q) + B) &= w \\ V'(q + Q(l_q))x(l_x)Q'(l_q) &= w. \end{aligned} \tag{7}$$

Entry We consider a continuum of competitive entrant firms that choose inventors l_e to produce innovations. At arrival rate $x_e(l_e)$, entrants take over a random product line, $j \in [0,1]$, improving the quality by a step of $Q_e > 0$. To concentrate on the speed-quality tradeoff of incumbents, we suppose quality Q_e is not a choice variable for the entrant firms.

The value function for entrant firms is,

$$rV_e - \dot{V}_e = \max_{l_e} \{x_e(l_e)[E[V(q_j + Q_e)] - V_e] - l_e w\} - \psi_e,$$

where the capital gains equal the expected value of innovation and market entry. Suppose the arrival rate is increasing, using the concave function. The FOC condition for entrants implies,

$$x_e'(l_e)[E[V(q_j + Q_e)] - V_e] = w. \tag{8}$$

If we assume free entry, entering firms exhaust profit opportunities. If $x_e(\cdot)$ has decreasing returns to scale, the adequate value for $\psi_e > 0$ guarantees that $V_e = 0$.

Growth Rate The quality of a product line evolves for a short time interval Δt by increasing stochastically according to the quality improvements chosen by the firm,

$$q_j(t + \Delta t) = q_j(t) + \begin{cases} Q(l_{qj}) & \text{with prob. } x(l_{xj})\Delta t \\ Q_e & \text{with prob. } x_e(l_e)\Delta t \\ 0 & \text{with prob. } 1 - (x_e(l_e) + x(l_{xj}))\Delta t \end{cases}$$

The average quality, therefore, grows at a rate,

$$g(t) = \frac{\dot{\bar{q}}(t)}{\bar{q}(t)} = \lim_{\Delta t \to 0} \frac{\bar{q}(t+\Delta t) - \bar{q}(t)}{\bar{q}(t)\Delta t} = \frac{\int_0^1 x(l_{xj}(t))Q(l_{qj}(t))dj + x_e(l_e(t))Q_e(t)}{\bar{q}(t)}. \tag{9}$$

For constant growth, the numerator in equation (9) must grow at the same rate as the aggregate quality of patents in the economy. Below, we discuss the conditions balanced growth imposes on the equilibrium.

Preferences and Market Clearing To close the model, we assume all workers have the same standard log preferences over the final good consumption $C(t)$ and characterize their aggregate decision using an infinitely lived representative household,

$$\int_0^\infty e^{-\rho t} \log C(t) dt \tag{10}$$

where ρ is the discount rate. Solving the maximization problem of the representative household yields the usual Euler equation,

$$\frac{\dot{C}(t)}{C(t)} = r(t) - \rho.$$

Workers are endowed with a unit of labor that they supply inelastically. Labor markets are competitive, with firms and workers taking prices as given. The labor market clearing condition is,

$$\int_0^1 (l_{xj} + l_{qj})dj + l_e = L_I. \tag{11}$$

Equilibrium The *competitive equilibrium* is characterized by a set of allocations $\{Y(t), C(t), k_j(t), l_{xj}(t), l_{qj}(t), l_e(t)\}$, prices $\{p_j(t), w(t), r(t)\}$ and the incumbent value function $V(q(t))$, such that: i) the final good producers maximize profits and aggregate output satisfies (5), ii) incumbent value function satisfies the HJB equation (6) and firms optimally choose labor for speed and quality solving (7), iii) entrant firms optimally choose labor for innovation solving (8), iv) the representative household solves the intertemporal maximization problem (10), and v) all labor and product markets clear.

Balanced Growth We study the economy on a balanced growth path (BGP), where the growth in average quality $\bar{q}(t)$, aggregate output $Y(t)$ and aggregate consumption $C(t)$ is constant. From equation (5), the growth in aggregate output is equal to the growth in the average quality. Moreover, from the final good market clearing, aggregate consumption equals aggregate output, so consumption also grows at the same rate, $g := \frac{\dot{\bar{q}}(t)}{\bar{q}(t)} = \frac{\dot{Y}(t)}{Y(t)} = \frac{\dot{C}(t)}{C(t)}$. As usual, this implies that the interest rate is constant on the BGP, $r(t) = r$.

Parametrization Suppose a general parametrization of the functional forms, where the speed $x(l_x)$, quality $Q(l_q)$ and additional benefit B are functions of the quality of their product line q and the average quality of the economy $\bar{q}$,

$$x = \chi l_x^{\alpha_x} q^{\gamma_x} \bar{q}^{\bar{\gamma}_x} \ , \ \ Q = \lambda l_q^{\alpha_q} q^{\gamma_q} \bar{q}^{\bar{\gamma}_q} \ , \ \ B = \chi_b q^{\gamma_b} \bar{q}^{\bar{\gamma}_b}.$$

We assume speed and quality have constant elasticities of labor, $\alpha_x > 0$ and $\alpha_q > 0$. The parameters γ_x, γ_q, and γ_b capture the elasticity of speed, quality, and benefit to the quality of the product line. Similarly, γ_x, γ_q, and γ_b reflect these elasticities with respect to the average quality of products in the economy. This is a flexible formulation that allows the speed or quality of innovations to increase or decrease as firms grow. Below, we show that balanced growth implies some restrictions on these parameters.

Consistent with these parametric assumptions, we postulate the value function is linear on the product line quality, $V(q_j) = Aq_j$. Replacing in the FOC,

$$\begin{aligned} [l_x]:\ & \alpha_x \chi l_x^{\alpha_x - 1} q^{\gamma_x} \bar{q}^{\bar{\gamma}_x} \left(A\lambda l_q^{\alpha_q} q^{\gamma_q} \bar{q}^{\bar{\gamma}_q} + \chi_b q^{\gamma_b} \bar{q}^{\bar{\gamma}_b} \right) = w \\ [l_q]:\ & \alpha_q \chi l_x^{\alpha_x} q^{\gamma_x} \bar{q}^{\bar{\gamma}_x} A\lambda l_q^{\alpha_q - 1} q^{\gamma_q} \bar{q}^{\bar{\gamma}_q} = w. \end{aligned} \tag{12}$$

For the value function to be linear in q, labor demands must be linear in $\frac{q}{\bar{q}}$ and $\frac{w}{\bar{q}}$ to be constant. The latter condition implies that wages must grow at the same rate as the average quality in the economy. Moreover, we study the economy on a BGP. Together, these conditions imply,

$$\gamma_x + \gamma_q = 1 - \alpha_x - \alpha_q,\ \bar{\gamma}_x + \bar{\gamma}_q = \alpha_x + \alpha_q,\ \gamma_x + \gamma_b = 1 - \alpha_x,\ \gamma_x + \bar{\gamma}_b = \alpha_x.$$

Combining these conditions implies the elasticity of the additional benefit with respect to quality, and the average quality of the economy must be equal to the sum of the labor elasticity and quality elasticities for the quality function, $\gamma_b = \alpha_q + \gamma_q$ and $\bar{\gamma}_b = \alpha_q + \bar{\gamma}_q$.

Lemma 1. *Suppose $\alpha_x + \alpha_q < 1$, then labor demands for speed and quality are given by $l_x = C_x(A)\frac{q}{\bar{q}}$ and $l_q = C_q(A)\frac{q}{\bar{q}}$, where $C_x(A)$ and $C_q(A)$ solve,*

$$C_x^{1-\alpha_x} = \frac{\alpha_x \bar{q}}{w} \left((\chi A \lambda)^{\frac{1}{1-\alpha_q}} \left(\frac{\alpha_q \bar{q}}{w} \right)^{\frac{\alpha_q}{1-\alpha_q}} C_x^{\frac{\alpha_x \alpha_q}{1-\alpha_q}} + \chi \chi_b \right),\ C_q = \left(\frac{\alpha_q \chi A \lambda \bar{q}}{w} \right)^{\frac{1}{1-\alpha_q}} C_x^{\frac{\alpha_x}{1-\alpha_q}}.$$

The condition $\alpha_x + \alpha_q < 1$ in Lemma (1) guarantees unique solutions for $C_x(A)$ and $C_q(A)$, for any the wage rate $w > 0$. Replacing the labor demands in the HJB equation for incumbent firms implies,

$$(r+\tau)A = \pi_I + \pi_q(A) + \pi_b(A),$$

where $\pi_q(A) = \chi A \lambda C_x(A)^{\alpha_x} C_q(A)^{\alpha_q}(1-\alpha_x-\alpha_q)$, $\pi_b(A) = \chi\chi_b C_x(A)^{\alpha_x}(1-\alpha_x)$.

For entrant firms, we suppose that the arrival rate and the quality of innovations for a product line with quality q are,

$$x_e = \chi_e l_e^{\alpha_e}, \; Q_e = \lambda_e(\nu q + (1-\nu)\bar{q}),$$

where α_e denotes the labor elasticity of the arrival rate of entrant innovations and $\nu \in [0,1]$ measures the extent to which the entrant innovation is built upon the quality of the product line on which it innovates. If $\nu \neq 1$, the average quality in the economy affects the step size of the entrant innovations. This can be interpreted as knowledge spillovers and implies a non-degenerate invariant distribution $F(\hat{q})$ for the normalized qualities in the economy $\hat{q} = \frac{q}{\bar{q}}$.

Since we suppose that entrant firms randomly innovate over the continuum of product lines, their labor demand is,

$$l_e = \left(\frac{\alpha_e \chi_e A(1+\lambda_e)\bar{q}}{w}\right)^{\frac{1}{1-\alpha_e}} =: C_e.$$

The market clearing condition in this case is,

$$\int \left(l_x(\hat{q}) + l_q(\hat{q})\right) dF(\hat{q}) + l_e = C_x + C_q + C_e = L_I.$$

Finally, we can compute the growth rate in the economy that depends on the speed and quality of incumbent innovation and the entrant firm innovation,

$$g = \frac{\int xQ + x_e(Q_e - q)}{\bar{q}} dF(\hat{q}) = \chi\lambda C_x^{\alpha_x} C_q^{\alpha_q} + \chi_e C_e^{\alpha_e}\lambda_e.$$

Welfare comparisons For our counterfactual exercises, we compare welfare in consumption equivalent units. More precisely, on a balanced growth path, consumption grows at a constant rate g, so $C(t) = C_0 e^{gt}$. The aggregate discounted utility is,

$$U(C_0, g) = \int_0^\infty e^{-\rho t} \log C_t dt = \int_0^\infty e^{-\rho t}\left(\log C_0 + gt\right) dt = \frac{g + \rho \log C_0}{\rho^2}.$$

This means that welfare rises if the initial level of goods produced and consumed increases or growth increases. Later in our quantitative exercises, we compare the welfare between the benchmark economy (C_0^0, g^0) and a counterfactual economy (C_0^c, g^c) by considering the consumption equivalent changes in welfare ω: $U(\omega C_0^0, g^0) = U(C_0^c, g^c) \implies \omega = e^{\frac{g^c - g^0}{\rho}} \frac{C_0^c}{C_0^0}$. In the case where initial consumption is the same across models, the change in welfare only stems from differences in the growth rate.

4.2 Application: Rising Concentration of Inventors in Large Firms

Recent studies have shown that economic activity and innovation have become increasingly concentrated in large firms, linked to a slowdown in growth and productivity (Akcigit and Ates 2021). This trend is also reflected in the concentration of inventors. Figure 6 illustrates the rise in human capital concentration in large firms using US patent data. The figure shows the share of inventors in the largest 10% of firms rose sharply from approximately 72% in 1980 to 81% in 2015.

Figure 6: Inventor Concentration

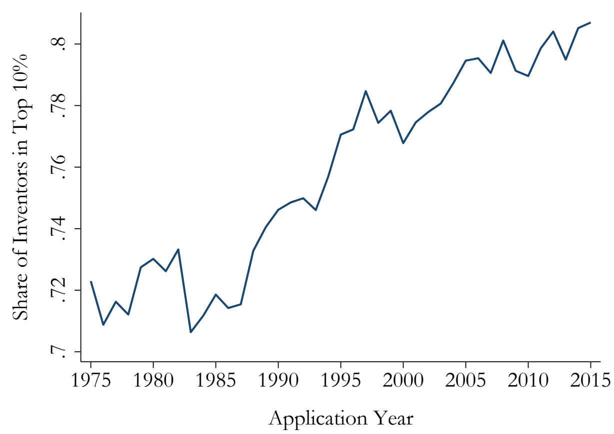

Notes: Share of inventors in top decile patenting firms by year.

The increasing inventor concentration in large firms could have different implications depending on the firm innovation strategies. For example, if large firms produce faster patents but of lower quality, this could negatively impact growth. Our objective is to use our estimated model to quantify the impact of the rising concentration on the allocation of inventors across the speed-quality margins and on growth.

Calibration We estimate the model using data on inventors and firm patents, aiming to match the key empirical moments. Specifically, we match the moments for the 10% largest firms and the 90% smallest firms based on the facts established in our analysis.

We start by calibrating the quantitative version of the model with the parameter values provided in Table 2 for two periods: 1980-1985 and 2010-2015. Specifically, we externally calibrate and normalize the supply of inventors L_I, the initial innovation productivity of incumbent firms χ in 1980, the initial private benefit parameter χ_b in 1980, the discount rate ρ, the elasticity of labor in final good production β, and the supply of production workers L. We choose the elasticity of innovation to speed and quality α_x and α_q to match the coefficients of the regression of log inventors on log patents and log quality for the period 1980-1985. The remaining parameters are estimated using the simulated method of moments.

Table 2: Parameters

Parameter	Description	Value 1980	Value 2010
Estimated using SMM			
–Levels (Both Periods)–			
λ	Scale Parameter of Quality	2.078	1.09
ν	Knowledge Diffusion for Entry	0.134	0.484
χ_e	Productivity Entrants	0.009	0.008
λ_e	Quality Step-Size Entrants	0.83	0.588
α_e	Entrant Innovation Elasticity of Labor	0.252	0.277
γ_q	Elasticity of Incumbent Quality	-0.085	-0.105
–Changes (Only 2010)–			
χ	Productivity Incumbents	0.1	0.13
χ_b	Additional Benefit	0	1.9
α_x	Speed Incumbent Innovation Elasticity of Labor	0.4	0.435
α_q	Quality Incumbent Innovation Elasticity of Labor	0.1	0.1
Externally Calibrated and Normalized Parameters			
L_I	Supply of Inventors	0.01	0.02
ρ	Discount rate	0.02	0.02
β	Elasticity of Labor of Production Workers	0.11	0.11
L_p	Supply of Production Labor	1	1

Notes: Estimated parameters for the period 1980-1985 and 2010-2015.

To estimate the parameters, we target ten crucial data moments in the two time periods: 1980-1985 and 2010-2015. Given the normalization and externally calibrated parameters in 1980-1985, we estimate 6 parameters targeting 7 moments in levels. These moments include the average growth rate from Garcia-Macia et al. (2019), the share of inventors in entrant firms, the share of entrant innovations, the relative quality of patents, the concentration of inventors at the top 10 firms, the concentration of patenting at top 10 firms and the patent quality at top 10 firms relative to bottom 90. Additionally, we take 1980-1985 as the base period fixing the value of χ and χ_b and target the change in the patenting speed, the change in the ratio of private value to patent quality, and the change in inventors per patent. The fit of the estimated model is presented in Table 3, demonstrating a close alignment between the model and the targeted moments.

Table 3: Moments Fit

Moment	Data 1980	Model 1980	Data 2010	Model 2010
–Levels (Both Periods)–				
Average Growth Rate	1.66%	1.66%	1.32%	1.32%
Inventors in Entrant Firms	8.2%	8.2%	4%	4%
Share of Entrant Innovations	9.7%	9.7%	5%	5%
Patent Quality: Q_{ent}/Q_{inc}	0.841	0.841	0.951	0.951
Inventors in Top 10	72.1%	70.4%	80%	79.5%
Share of Innovations in Top 10	68.7%	68.8%	77.8%	79.1%
Patent Quality: Top 10 Relative to Bottom 90	1.169	1.17	0.924	0.932
–Changes (Only 2010)–				
Change in Speed: $(x/x_{base\ year})$	-	-	1.56	1.59
Change in Private Value/Quality: $(\log{(V/Q)})$	-	-	2.43	2.61
Change in Inventors per Patent	-	-	1.47	1.37

Notes: Data targeted moments and model fit for the period 1980-1985 and 2010-2015.

The targeted moments exhibit revealing patterns across the two periods. Despite a notable 56% rise in patenting speed and a doubling of the total number of inventors, the targeted growth decreases from 1.66% to 1.32%, in line with the productivity slowdown documented in recent literature (Bloom et al., 2020, Akcigit and Ates, 2021, and Aghion et al., 2023). The parameters governing the innovation production function must adjust to rationalize these changes. Specifically, the productivity of incumbent innovation χ rose by 30% from 0.1 to 0.13, while the scale parameter λ declined by almost 50%. Moreover, there is a diminishing importance of entrant innovation and a concentration of patenting among the top 10 firms. The reduced scope for entrants entails a 10% decrease in their productivity parameter χ_e, with entrant innovations increasingly reliant on product line quality ν, rising from 0.13 to 0.48, echoing the knowledge diffusion decline highlighted by Akcigit and Ates (2021). The speed-quality tradeoff is apparent from the observed changes in relative patent quality. Despite faster patenting by incumbent firms, entrants' patent quality relative to incumbents increased from 0.84 to 0.95, while the top 10 firms' patent quality relative to the bottom 90 declined from 1.17 to 0.92. Finally, there is a substantial surge in the private value relative to patent quality, increasing by a factor of 2.4 between the two periods. To capture this change, the additional private benefit parameter χ_b rises to 1.9, equating to 20% of total profits from the private benefit. Collectively, these parameter adjustments explain the observed changes in speed and quality of innovation, which we further decomposed in the subsequent analysis to understand the consequences on growth and the allocation of inventors.

Growth Decomposition and the Allocation of Inventors Using the estimated model, we can perform a decomposition analysis to discern the contributions of speed and quality components to the observed slowdown in growth. In Table 4 we decompose growth into percentage changes, where

$$\Delta\%g = \underbrace{\left(\frac{x}{x_e + x}\Delta\%x + \frac{x_e}{x_e + x}\Delta\%x_e\right)}_{x^{tot}:\ Speed} + \underbrace{\left(\frac{Q}{Q_e + Q}\Delta\%Q + \frac{Q_e}{Q_e + Q}\Delta\%Q_e\right)}_{Q^{tot}:\ Quality} + \underbrace{\Delta\%x^{tot}\Delta\%Q^{tot}}_{Speed \times Quality}.$$

Consistent with prior research (Garcia-Macia et al., 2019), our decomposition highlights the predominant contribution of incumbents to overall growth. Over time, the importance of incumbents has grown, with their share of total growth reaching 95% of total growth, compared to 93% in the 1980s. The observed decline in targeted growth can be attributed to reduced growth rates for both incumbents and entrants. However, while entrants experienced a notable decrease in the number of innovations, the negative impact of incumbents on growth arises from the decline in quality despite a pronounced increase in speed during the period.

Table 4: Growth Decomposition

	Growth Rate		**Percentage Change**			
Firms	*1980*	*2010*	*Speed*	*Quality*	*Speed × Quality*	*Total*
Incumbent	1.54%	1.25%	53.4%	-46.9%	-25%	-18.5%
Entrant	0.12%	0.07%	-25.7%	-29.2%	7.5%	-47.4%
Total	1.66%	1.32%	45.7%	-39.3%	-27.1%	-20.6%

Notes: Growth rate decomposition into percentage changes for incumbents and entrants.

Table 5 presents the distribution of inventors among entrant firms and across speed and quality categories for incumbent firms. The table reveals that a significant portion of the decrease in inventors allocated to entrant innovations is redirected toward increasing the speed of incumbent innovations, which grows from 74% to 86% of total labor. This reallocation results in a decline in the proportion of labor dedicated to quality from 28% to 10%, significantly impacting the overall quality of incumbent innovations.

Table 5: Allocation of Inventors

Inventors	**1980**	**2010**
Speed (Incumbents)	73.5%	85.7%
Quality (Incumbents)	18.3%	10.3%
Entrant Firms	8.2%	4%

Notes: Allocation of inventors across speed and quality in 1980-1985 and 2010-2015.

We can further decompose incumbent speed and quality into the change of variables and parameters as follows,

$$\Delta \log x = \Delta \log(\chi) + \alpha_x \Delta \log(l_x) + \Delta \alpha_x \log(l_x)$$
$$\Delta \log Q = \Delta \log(\lambda) + \alpha_q \Delta \log(l_q) + \Delta \alpha_q \log(l_q),$$

where Δ denotes the change between the two periods.

Table 6 shows that an increase in productivity and labor drives the increase in speed. The productivity is larger, labor doubles, and there is reallocation towards speed. In contrast, the deterioration in quality stems from less productivity and reallocation of labor away from quality.

Table 6: Speed and Quality Decomposition

	Productivity $\Delta \log(\chi) \vee \Delta \log(\lambda)$	Labor Quantity $\alpha_i \Delta \log(l_{inc})$	Labor Allocation $\alpha_i \Delta \log(l_i/l_{inc})$	Elasticity $\Delta \alpha_i \log(l_i)$	Total
Speed	0.223	0.296	0.044	-0.164	0.428
Quality	-0.645	0.074	-0.062	0.001	-0.632

Notes: Speed and quality decomposition into changes in productivity, labor and elasticities.

Lastly, in Figure 7, we explore the effect of changing the proportion of labor allocated to the speed and quality on the arrival rate, the quality of innovations, and growth. The red dashed lines represent the estimated parameter values for the period 1980-1985 (the diamond denotes the calibration point), while the solid blue line corresponds to the estimated parameters for 2010-2015 (the circle indicates the calibration point). Intuitively, Figures 7a and 7b illustrate that as labor allocated to speed increases, the arrival rate rises while quality declines. Over time, speed has become less concave, pulling more labor towards speed. Figure 7c indicates that most of the decline in growth stems from changes in innovation production, reflected in the downward shift of the curve between the two periods. Even though there has been a significant increase in labor allocated to speed, the changes in innovation production technology attenuate the effects on growth. Figure 7d shows that reallocating labor across speed and quality implies a small increase the growth rate from 1.32% to 1.34%, with around 80% of labor still allocated to speed. However, if the proportion of labor allocated to speed in 2010-2015 is imposed to the 1980-1985 calibration, growth would be more significantly impacted, going down from 1.66% to 1.59%.

Figure 7: Counterfactual– Changing Speed/Quality Labor

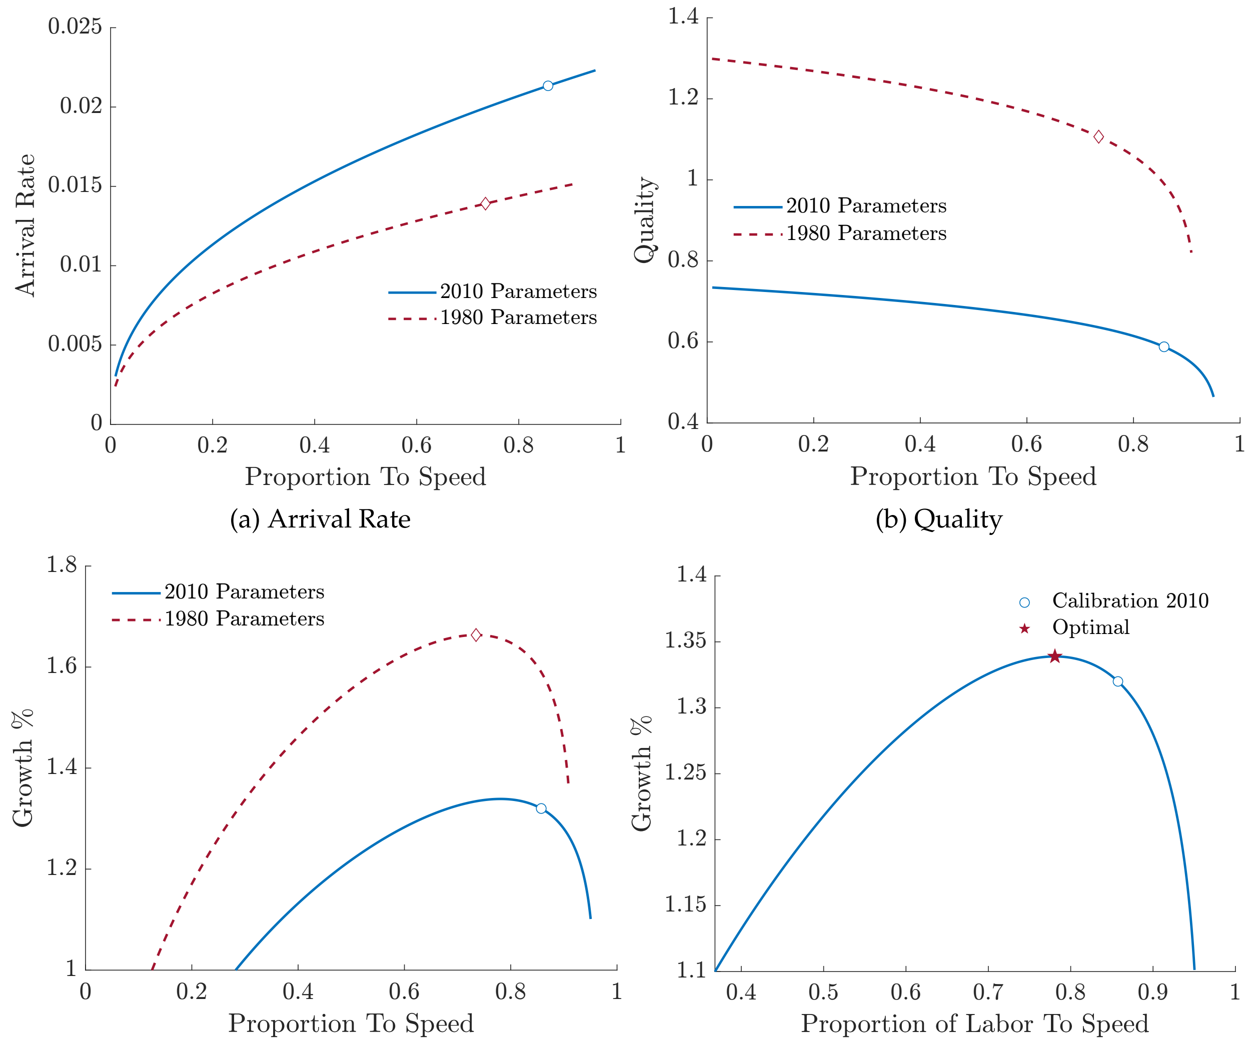

(a) Arrival Rate (b) Quality (c) Growth (d) Optimal Growth

Notes: Counterfactual changing the proportion of labor going to speed keeping labor for entrants fixed, using 1980-1985 parameters (dotted line) and 2010-2015 parameters (solid line). Calibration points are depicted as a diamond (1980-1985) and a circle (2010-2015).

Endogenous Quality Even in the case of modest effects of the reallocation of labor across speed and quality, endogenizing quality is still important. To illustrate this point, we consider an economy with the same parameters as the estimated model but with fixed quality. To compare both scenarios, we choose the scale parameter of quality λ in the counterfactual economy to match the observed growth for the period.

Figure 8: Counterfactual–Endogenous vs. Fixed Quality (1980-1985)

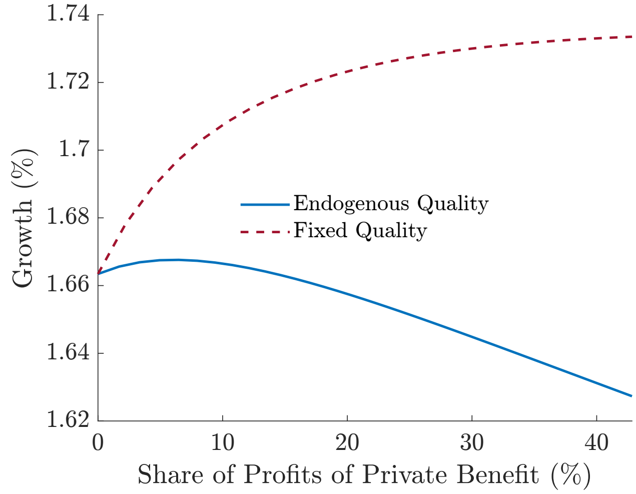

Notes: Counterfactual comparing model with endogenous quality (solid line) and fixed quality (dotted line) using 1980-1985 parameters.

Figure 8 illustrates the impact of increasing the additional private benefit parameter χ_b on the growth rate in two scenarios: one with fixed quality and the other with endogenous quality, calibrated for 1980-1985. Initially set to $\chi_b = 0$, both models start from the targeted growth rate. To ease interpretation, we translate changes in χ_b into the share of profits derived from the private benefit. Notably, the models exhibit contrasting trajectories. Initially, a larger private benefit squeezes entry in both models, resulting in larger growth. As entrants produce a negative externality from the business-stealing effect, in equilibrium, they command excessive labor. However, while growth keeps increasing with fixed quality (depicted by the red dotted line), a larger private benefit in the endogenous quality model lowers quality, thereby suppressing growth. The difference in the growth rate between the two models is substantial as the private benefit growth as a share of total profits.

More broadly, in models where entry isn't excessive, speed creates another avenue for potential misallocation. Increasing private benefits might decrease growth by diminishing quality and constraining entry. Our ongoing work aims to calibrate an extended model that incorporates and quantifies this additional channel.

5 Conclusion

In this paper, we study the interplay between the speed and quality of innovations and their importance for growth. We consider the speed-quality tradeoff at the center of an endogenous growth model informed by micro-data on inventors. We show large firms tend to pursue innovative strategies that use inventors in faster but less socially valuable projects. This creates a potential misallocation in the aggregate economy, which comes through the links between concentration, talent allocation across firms, innovation, and economic growth.

Our exploration of the *speed-quality tradeoff* underscores its pivotal role in shaping innovation outcomes. Leveraging both empirical evidence and our quantitative exercises, we analyze the nuanced effects of reallocating labor between speed and quality on innovation and growth. While our findings reveal modest improvements in growth through optimal labor allocation, they also underscore the crucial role of endogenizing quality in understanding innovation dynamics.

Our ongoing analysis seeks to analyze the scenarios where the speed and quality tradeoff matter the most, for instance, by explicitly considering the influence of private benefits on firm entry and innovation outcomes.

References

Abrams, D. S.: 2008, 'Did Trips Spur Innovation-An Analysis of Patent Duration and Incentives to Innovate'. *University of Pennsylvania Law Review* **157**, 1613.

Acemoglu, D.: 2008, *Introduction to modern economic growth*. Princeton university press.

Acemoglu, D. and U. Akcigit: 2012, 'Intellectual property rights policy, competition and innovation'. *Journal of the European Economic Association* **10**(1), 1–42.

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

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li: 2023, 'A Theory of Falling Growth and Rising Rents'. *The Review of Economic Studies*. rdad016.

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

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

Akcigit, U. and S. Ates: 2023, 'What Happened to U.S. Business Dynamism?'. *Journal of Political Economy* (Forthcoming).

Akcigit, U. and S. T. Ates: 2021, 'Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory'. *American Economic Journal: Macroeconomics* **13**(1), 257–98.

Akcigit, U., S. Caicedo, E. Miguelez, S. Stantcheva, and V. Sterzi: 2018, 'Dancing with the Stars: Innovation through Interactions'. *National Bureau of Economic Research Working Paper # 24466*.

Akcigit, U., J. Grigsby, and T. Nicholas: 2017, 'The Rise of American Ingenuity: Innovation and Inventors of the Golden Age'. *National Bureau of Economic Research Working Paper # 23047*.

Akcigit, U. and W. R. Kerr: 2018, 'Growth through Heterogeneous Innovations'. *Journal of Political Economy* **126**(4), 1374–1443.

Akcigit, U., J. G. Pearce, and M. Prato: 2020, 'Tapping into Talent: Coupling Education and Innovation Policies for Economic Growth'. *National Bureau of Economic Research Working Paper # 27862*.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen: 2020, 'The Fall of the Labor Share and the Rise of Superstar Firms*'. *The Quarterly Journal of Economics* **135**(2), 645–709.

Azoulay, P., J. S. Graff Zivin, and J. Wang: 2010, 'Superstar Extinction'. *Quarterly Journal of Economics* **125**(2), 549–589.

Becker, G. S.: 1960, 'An Economic Analysis of Fertility'. In: *Demographic and economic change in developed countries*. Columbia University Press, pp. 209–240.

Becker, G. S. and H. G. Lewis: 1973, 'On the interaction between the quantity and quality of children'. *Journal of Political Economy* **81**(2, Part 2), S279–S288.

Bertolotti, F.: 2022, 'Patent Length, Innovation, and the Role of Technology Disclosure Externalities'.

Bessen, J. E.: 2017, 'Information Technology and Industry Concentration'. Law & Economics Paper Series 17–41, Boston University School of Law.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb: 2020, 'Are ideas getting harder to find?'. *American Economic Review* **110**(4), 1104–44.

Bloom, N., R. Sadun, and J. Van Reenen: 2016, 'Management as a Technology?'. Technical report, National Bureau of Economic Research.

Covarrubias, M., G. Gutiérrez, and T. Philippon: 2020, 'From Good to Bad Concentration? US Industries over the Past 30 Years'. *NBER Macroeconomics Annual* **34**, 1–46.

Crouzet, N. and J. Eberly: 2018, 'Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles'. Jackson Hole Economic Policy Symposium.

Cunningham, C., L. Foster, C. Grim, J. Haltiwanger, S. W. Pabilonia, J. Stewart, and Z. Wolf: 2021, 'Productivity dispersion, entry, and growth in us manufacturing industries'. Technical report, Bureau of Labor Statistics.

De Loecker, J., J. Eeckhout, and G. Unger: 2020, 'The Rise of Market Power and the Macroeconomic Implications*'. *Quarterly Journal of Economics* **135**(2), 561–644.

De Ridder, M.: 2019, 'Market Power and Innovation in the Intangible Economy'. Cambridge Working Papers in Economics 1931, Faculty of Economics, University of Cambridge.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda: 2016, 'Declining business dynamism: What we know and the way forward'. *American Economic Review* **106**(5), 203–207.

Eggertsson, G. B., J. A. Robbins, and E. G. Wold: 2018, 'Kaldor and Piketty's Facts: The Rise of Monopoly Power in the United States'. National Bureau of Economic Research Working Paper 24287.

Garcia-Macia, D., C.-T. Hsieh, and P. J. Klenow: 2019, 'How destructive is innovation?'. *Econometrica* **87**(5), 1507–1541.

Garicano, L. and E. Rossi-Hansberg: 2006, 'Organization and Inequality in a Knowledge Economy*'. *Quarterly Journal of Economics* **121**(4), 1383–1435.

Grossman, G. M. and E. Helpman: 1991, 'Quality Ladders in the Theory of Growth'. *Review of Economic Studies* **29**(4), 43–61.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg: 2001, 'The NBER patent citation data file: Lessons, insights and methodological tools'.

Hill, R. and C. Stein: 2022, 'Race to the Bottom: Competition and Quality in Science'. *Working Paper*.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow: 2019, 'The Allocation of Talent and U.S. Economic Growth'. *Econometrica* **87**(5), 1439–1474.

Jaimovich, N. and S. Rebelo: 2017, 'Nonlinear Effects of Taxation on Growth'. *Journal of Political Economy* **125**(1), 265–291.

Jaravel, X., N. Petkova, and A. Bell: 2018, 'Team-Specific Capital and Innovation'. *American Economic Review* **108**, 1034–1073.

Jones, C. I.: 2023, 'Recipes and economic growth: A combinatorial march down an exponential tail'. *Journal of Political Economy* **131**(8), 000–000.

Kalyani, A.: 2022, ‘The Creativity Decline: Evidence from US Patents’. SSRN Working Paper.

Klenow, P. J. and H. Li: 2021, ‘Innovative Growth Accounting’. *NBER Macroeconomics Annual* **35**, 245–295.

Klette, T. and S. Kortum: 2004, ‘Innovating Firms and Aggregate Innovation’. *Journal of Political Economy* **112**(5), 986–1018.

Kline, P., N. Petkova, H. Williams, and O. Zidar: 2019, ‘Who Profits from Patents? Rent-Sharing at Innovative Firms*’. *Quarterly Journal of Economics* **134**(3), 1343–1404.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman: 2017, ‘Technological Innovation, Resource Allocation, and Growth’. *Quarterly Journal of Economics* **132**(2), 665–712.

Lentz, R. and D. T. Mortensen: 2008, ‘An empirical model of growth through product innovation’. *Econometrica* **76**(6), 1317–1373.

Liu, E., A. Mian, and A. Sufi: 2022, ‘Low Interest Rates, Market Power, and Productivity Growth’. *Econometrica* **90**(1), 193–221.

Lucas, R. E.: 1988, ‘On the Mechanics of Economic Development’. *Journal of Monetary Economics* **22**(1), 3 – 42.

Lucas, R. E. and B. Moll: 2014, ‘Knowledge Growth and the Allocation of Time’. *Journal of Political Economy* **122**(1), 1–51.

Ma, Y.: 2022, ‘Specialization in a knowledge economy’. *Available at SSRN 4052990*.

Manera, A.: 2022, ‘Competing for Inventors: Market Concentration and the Misallocation of Innovative Talent’. *Available at SSRN 4212815*.

Mincer, J.: 1984, ‘Human capital and economic growth’. *Economics of Education Review* **3**(3), 195 – 205.

Murphy, K. M., A. Shleifer, and R. W. Vishny: 1991, ‘The Allocation of Talent: Implications for Growth’. *Quarterly Journal of Economics* **106**(2), 503–530.

Pearce, J.: 2020, ‘Idea Production and Team Structure’. *Working Paper*.

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

USPTO: 2019, ‘PatentsView’. https://www.patentsview.org/download/. Updated December 31, 2019.

Waldinger, F.: 2016, 'Bombs, Brains, and Science: The Role of Human and Physical Capital for the Creation of Scientific Knowledge'. *Review of Economics and Statistics* **98**(5), 811–831.

Willis, R. J. and S. Rosen: 1979, 'Education and Self-Selection'. *Journal of Political Economy* **87**(5, Part 2), S7–S36.

Wu, L., D. Wang, and J. A. Evans: 2019, 'Large teams develop and small teams disrupt science and technology'. *Nature* **566**(7744), 378–382.

Appendix

A Empirical Appendix

The literature on innovation has proposed a panoply of measures to proxy innovations' social value and quality. In Figure A.1, we show the robustness of the macro trends in patent quality using the raw measure of Log 3-yr forward citations, not including self-citations, residualizing without team size control, considering 5-yr and 8-yr forward citations, Top 5% and Top 1% cited patents, originality and generality indexes.

- *Log 3-yr forward citations*: Raw measure of patent quality using the logarithm of one plus 3-yr forward citations.
- *Patent quality without self-citations*: Compute the 3-yr forward citations without including citation by the same assignee (firm). The measure is the residualized Log 3-yr citations (without self-citations) controlling for time, technology class (IPC3) and team size fixed effects.
- *Patent Quality without team size controls*: Residualized Log 3-yr citations controlling for time and technology class (IPC3).
- *Patent Quality using 5-yr and 8-yr forward citations*: Residualized Log 5-yr and 8-yr citations controlling for time, technology class (IPC3) and team size fixed effects.
- *Top 5% and Top 1% cited patents*: Compute the top 1% and top 5% among 3-yr forward citations and residualize measure controlling for time, technology class (IPC3) and team size fixed effects.

The last two measures are proposed by Hall et al. (2001). The Originality Index measures the dispersion of citations a patent gives,

$$Originality_j = 1 - \sum_{i \in I} s_{ij}^2,$$

where s_{ij} is the share of citations patent j makes technology class i.

The Generality Index measures the dispersion of citations a patent receives,

$$Generality_j = 1 - \sum_{i \in I} s_{ij}^2,$$

where s_{ij} is share of citations patent j receives from technology class i.

Figure [A.1](#) shows the time trend of various measures of patent quality and impact over time, split by the bottom 99% and top 1% firms.

Figure A.1: Measures of Patent Quality

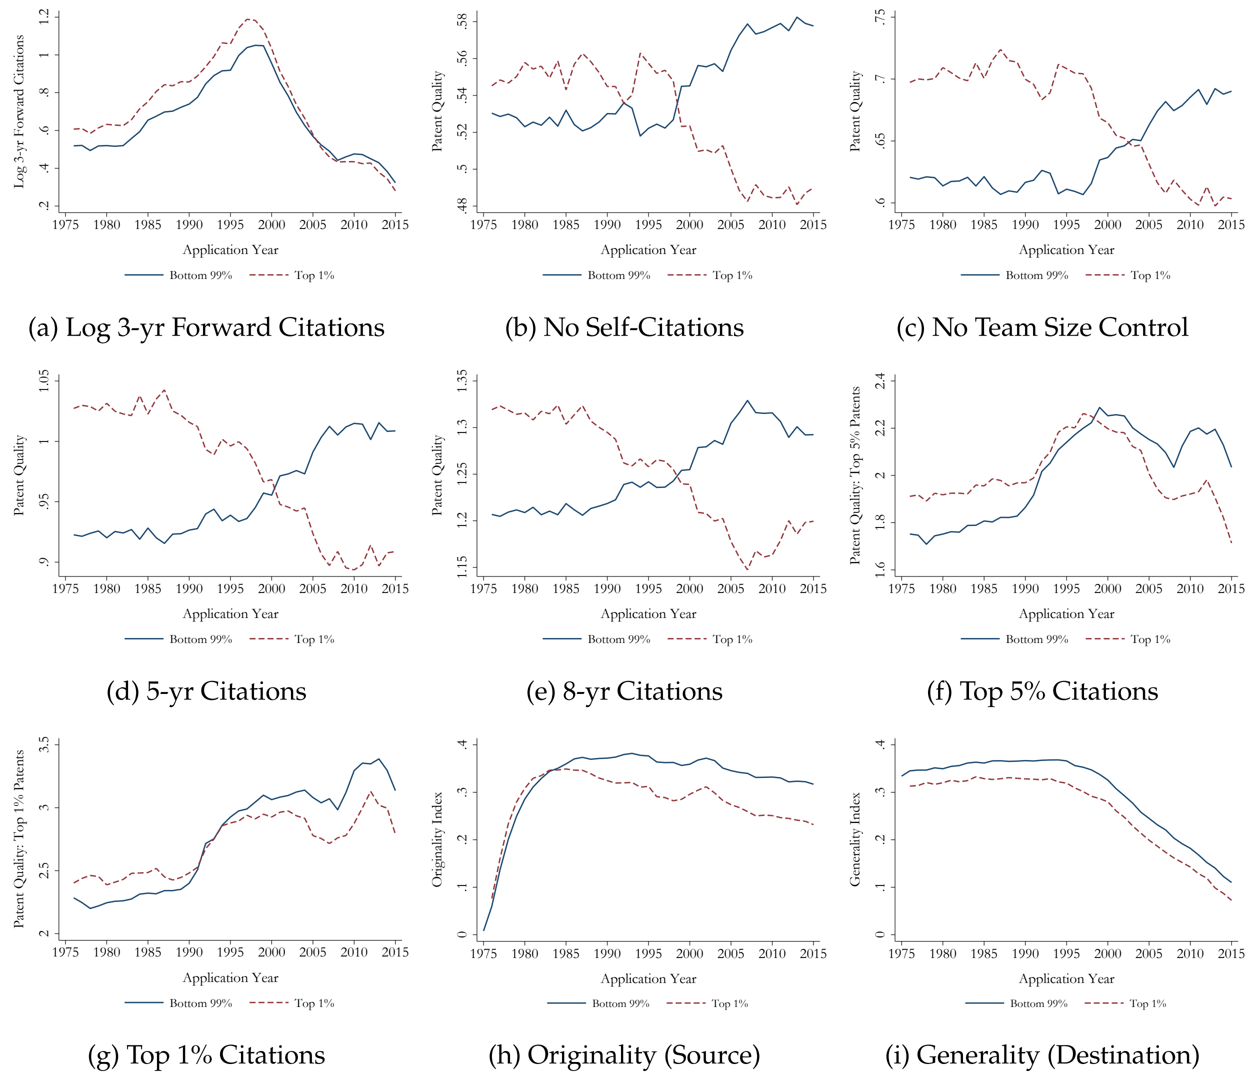

(a) Log 3-yr Forward Citations (b) No Self-Citations (c) No Team Size Control

(d) 5-yr Citations (e) 8-yr Citations (f) Top 5% Citations

(g) Top 1% Citations (h) Originality (Source) (i) Generality (Destination)

Notes: Patent quality using various measures over time. Refer to the text for the definition of each measure.

The measures in Figure A.1 tell a coherent story. There is a relative decline in quality at the top firms, as measured in a variety of ways.

Movers from Large to Small Firms

Figure A.2 studies the inventors who move in the opposite direction of Figure 3. These inventors move from a large firm to a small firm. The specification is the same as in Equation (3) We look at the same four outcome variables: arrival rate at the firm level, speed at the inventor level, patent quality, and patent private value.

Figure A.2: Inventor Moves From Large to Small Pre-1995 vs. Post-1995

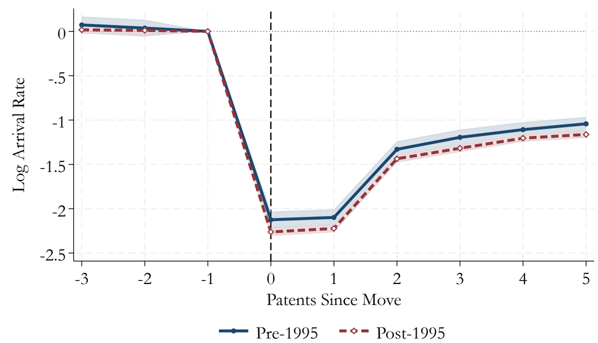

(a) Arrival Rate: Firm-Level

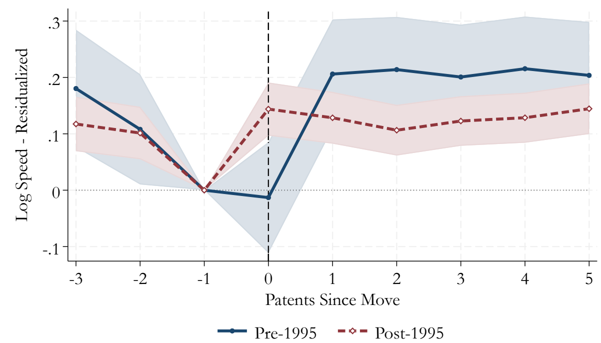

(b) Speed: At Inventor Level

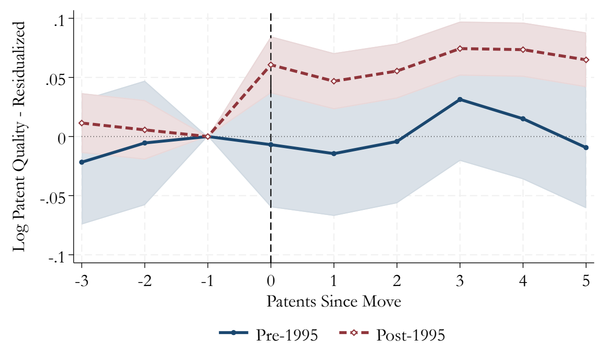

(c) Quality: Log Citations

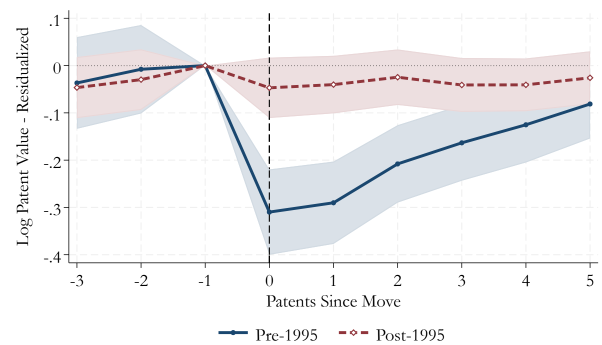

(d) Private Value: Log Stock Market Value

Notes: We measure the arrival rate at the firm level as the log number of patents per firm, the speed at the inventor level is the inverse of the average time to patent, patent quality=$Log(1 + 3yr\ forward\ citations)$, residualized using time, technology class, and team size controls and private value is the log of the stock market value change from Kogan et al. (2017). The regression model is specified in equation (3).

We see almost symmetric responses as in Figure A.2. The only difference is that in both cases movers experience an increase in speed around the time of the event, speaking to the fact that movers may speed up their innovations naturally. Beyond that, the private value declines upon moving, though this is less stark in the post period.

Notes on Natural Experiment

In the main text, we focused on the role of the policy shock in shifting firms' innovation decisions over speed and quality. In this section, we expand the discussion looking at splitting the data by other outcomes, firm size, and type of patenting.

Figure A.3 reports other outcome variables from the specification from Equation (4). In this

Figure A.3: Other Outcomes Policy Change

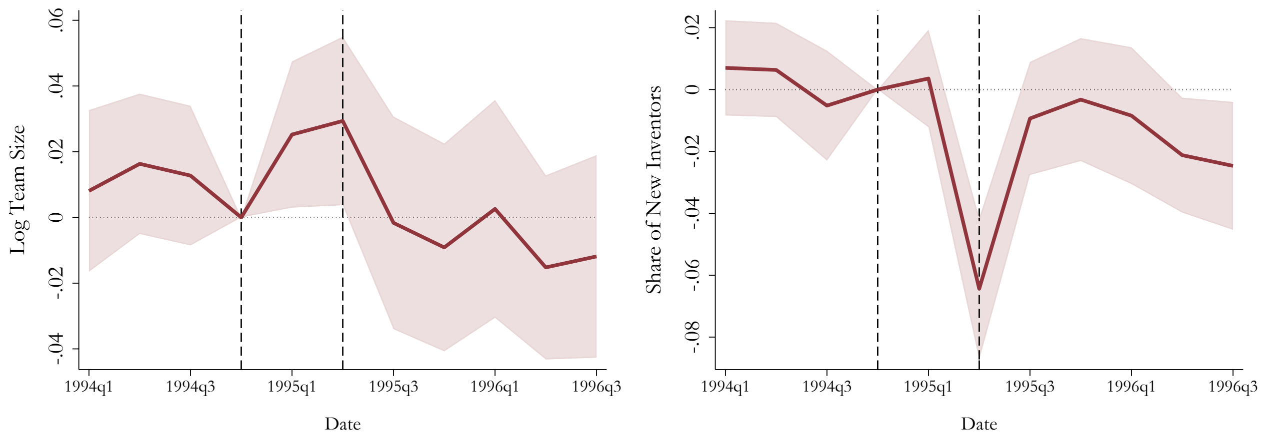

(a) Log Team Size (b) Entrant Inventors

Notes: Team size is the number of inventors in a patent. Entrant inventors are inventors that appear for the first time in the data. Plotted coefficients correspond to the interaction term of the Most Exposed (75th percentile) firms.

case, the dependent variable is team size and entrant inventors, which shows how firms substitute for more experienced inventors and larger teams around the time of the event.

We split by firm size in Figure A.4. Again, we look at the outcomes following the same specification as Equation (4).

Though large and small firms pursue very different innovation strategies on average, their response to the event is quite uniform. This is suggestive that the event can be used as a study of speed and quality across the firm size distribution. This uniformity is promising for the generalizability of the result.

Finally, in Figure A.5, we focus on the difference between internal and external patents. We again follow the specification in Equation 4 and split by internal and external patents. Unsurprisingly, firms are more capable of increasing their speed when it comes to internal patenting.

Figure A.4: Policy Change For Top 1% and Bottom 99%

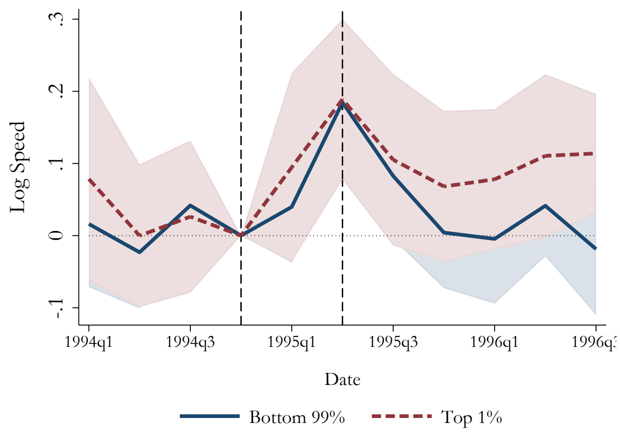

(a) Log Speed

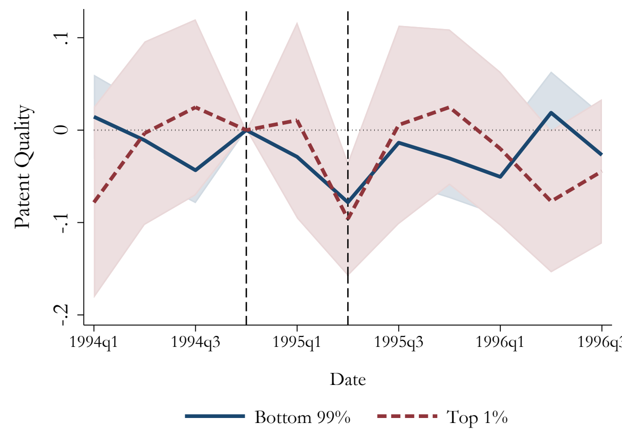

(b) Patent Quality

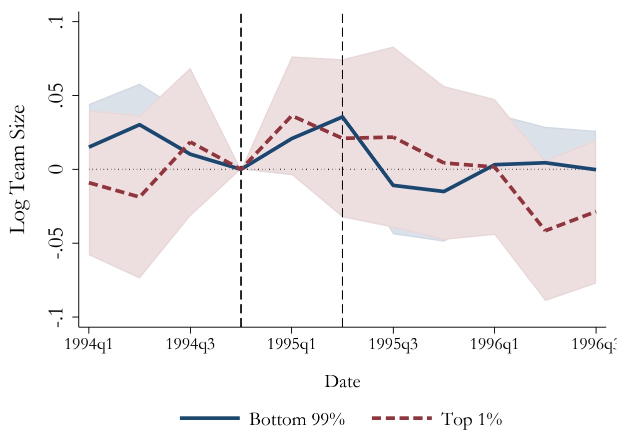

(c) Log Team Size

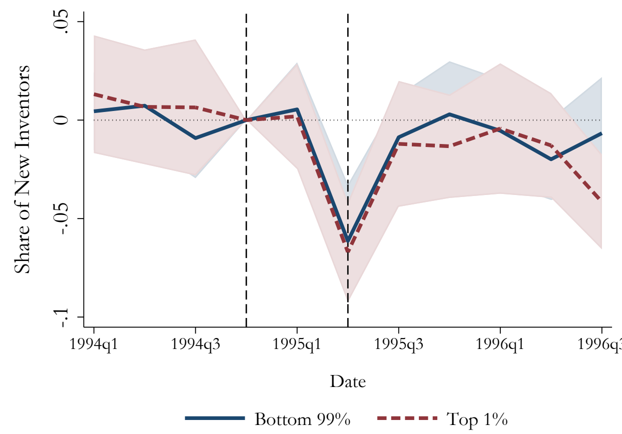

(d) Entrant Inventors

Notes: Speed is the inverse Time-to-Patent defined at the inventor level as the time it takes to produce the next patent, measured in days. Patent quality = $log(1 + 3yr\ fcit)$ residualized controlling for IPC3 technological class and year fixed effects. Team size is the number of inventors in a patent. Entrant inventors are inventors that appear for the first time in the data. Plotted coefficients correspond to the interaction term of the Most Exposed (75th percentile) firms. Firms split into Top 1% and Bottom 99% by number of patents by year.

Figure A.5: Policy Change For Internal and External Patents

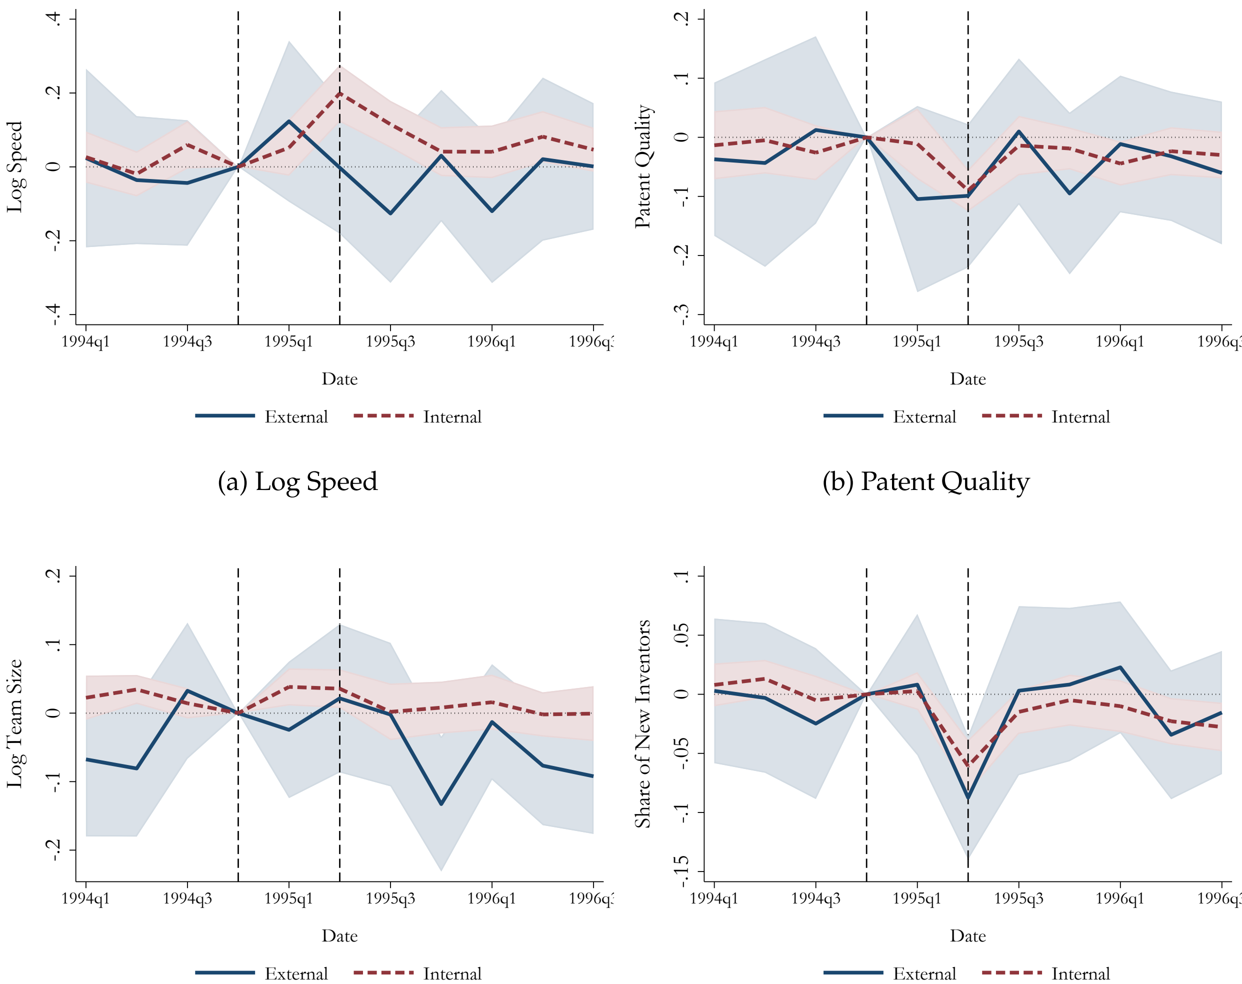

(a) Log Speed

(b) Patent Quality

(c) Log Team Size

(d) Entrant Inventors

Notes: Speed is the inverse Time-to-Patent defined at the inventor level as the time it takes to produce the next patent, measured in days. Patent quality = $log(1 + 3yr\ fcit)$ residualized controlling for IPC3 technological class and year fixed effects. Team size is the number of inventors in a patent. Entrant inventors are inventors that appear for the first time in the data. Plotted coefficients correspond to the interaction term of the Most Exposed (75th percentile) firms. Patents are split into external (no citations in previous IPC3 technological class) and internal (at least one citation in previous IPC3 technological class).

B Theoretical Appendix

B.1 Proofs

Proposition (Labor Allocation)**.** *The labor for speed relative to quality, $\frac{l_x^*}{l_q^*}$, is increasing in B.*

Proof. Given $w > 0$, $B \geq 0$ and the assumptions on x and Q, the first order conditions are also sufficient, and the solution implies demand functions l_x^*, l_q^* that satisfy,

$$[l_x]: \; x'(l_x^*)(V(q+Q(l_q^*)) - V(q) + B) = w$$
$$[l_q]: \; x(l_x^*)V'(q+Q(l_q^*))Q'(l_q^*) = w.$$

To simplify notation, we drop the arguments of the functions when there is no ambiguity and let $\Delta V := V(q+Q(l_q^*)) - V(q)$. We use the following lemma to do the comparative statics with respect to B.

Lemma. *$V(q)$ is increasing and weakly concave.*

Proof. $V(q)$ increasing follows from the instantaneous profit being increasing in q. For concavity given $\pi(q)$ is weakly concave, xQ is concave in (l_x, l_q) and $[0, L_I]$ is a convex set, then $V(q)$ is weakly concave (see Theorem 7.16 in Acemoglu 2008). □

The first-order conditions determine the optimal demand curves. Let $l_x^{[x]}$ denote the curve defined by equation $[l_x]$ and $l_q^{[q]}$ the curve with respect to $[l_q]$. By totally differentiating these curves, we can show that these mappings are increasing,

$$\frac{dl_x^{[x]}}{dl_q} = -\frac{x'V'Q'}{x''(\Delta V + B)} > 0 \; ; \; \frac{dl_q^{[q]}}{dl_x} = -\frac{x'V'Q'}{xV''(Q')^2 + xV'Q''} > 0, \quad \forall l_x, l_q \geq 0.$$

Moreover, taking the derivative again, we can show they are also concave,

$$\frac{d^2 l_x^{[x]}}{dl_q^2} = -\frac{x'x''\left(V'Q'' + x'V''(Q')^2(\Delta V + B) - (V'Q')^2\right)}{\left(x''(\Delta V + B)\right)^2} < 0,$$
$$\frac{d^2 l_q^{[q]}}{dl_x^2} = \frac{-x''V'Q'\left[xV''(Q')^2 + xV'Q''\right] + x'V'Q'\left[x'V''(Q')^2 + x'V'Q''\right]}{\left(xV''(Q')^2 + xV'Q''\right)^2} < 0.$$

Notice that for curve $[l_x]$, if $l_q = 0$, then $l_x^0 = (x')^{-1}(w/B)$, while curve $[l_q]$ starts at the origin. Figure B.1 illustrates the optimal demand curves.

Figure B.1: Labor Demand for Speed and Quality– Change in Private Value B

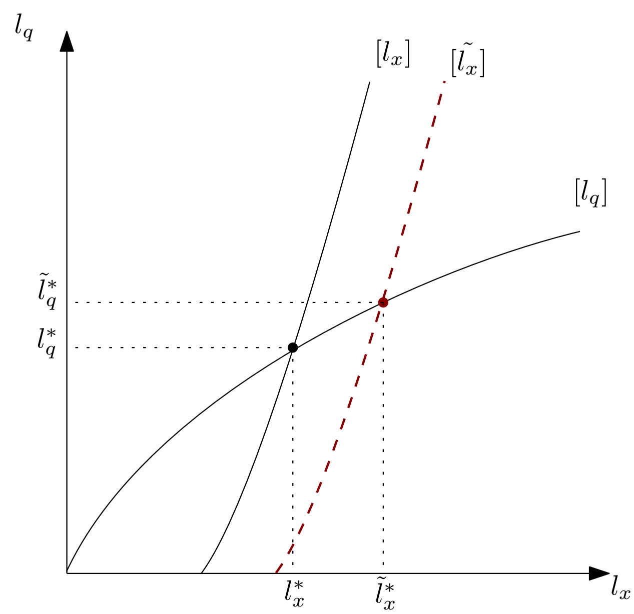

Now, using this characterization, we can do comparative statics with respect to B. If B increases to $\tilde{B} > B$, curve $[l_x]$ shifts to the right, while the $[l_q]$ does not change. Given that curve $[l_q]$ is concave, l_x increases relatively more than l_q. So l_x^*/l_q^* is increasing in B.

□

Proposition. *If $B > 0$, less relative labor to speed increases growth from innovating firms, xQ.*

Proof. The growth contribution of innovating firms is,

$$g_{inc} = x(l_x)Q(l_q).$$

Let $\tilde{L}_I > 0$ denote the mass of inventors to allocate across speed and quality.

The allocation that maximizes growth solves,

$$\max_{l_x, l_q} x(l_x)Q(l_q) \;\; s.t \;\; l_x + l_q = \tilde{L}_I.$$

The FOCs are,

$$x'(l_x)Q(l_q) = x(l_x)Q'(l_q)$$
$$l_x + l_q = \tilde{L}_I$$

Let l_x^g, l_q^g be the solution to the growth maximizing allocation.

We can compare the first condition to the labor allocation in equilibrium, where the marginal value of labor for speed and quality is equalized,

$$x'(l_x)(AQ(l_q) + B) = x(l_x)AQ'(l_q)$$

So if $B > 0$, the implicit function $l_x^*(l_q) \geq l_x^g(l_q)$, for all l_q with strict inequality for some l_q. Since $l_x + l_q = \tilde{L}$, then $l_x^*/l_q^* > l_x^g/l_q^g$.

$\square$

Quantitative Model

Lemma. *Suppose $\alpha_x + \alpha_q < 1$, then labor demands for speed and quality are given by $l_x = C_x(A)\frac{q}{\bar{q}}$ and $l_q = C_q(A)\frac{q}{\bar{q}}$, where $C_x(A)$ and $C_q(A)$ solve,*

$$C_x^{1-\alpha_x} = \frac{\alpha_x \bar{q}}{w}\left((\chi A\lambda)^{\frac{1}{1-\alpha_q}}\left(\frac{\alpha_q \bar{q}}{w}\right)^{\frac{\alpha_q}{1-\alpha_q}} C_x^{\frac{\alpha_x\alpha_q}{1-\alpha_q}} + \chi\chi_b\right),\ C_q = \left(\frac{\alpha_q\chi A\lambda\bar{q}}{w}\right)^{\frac{1}{1-\alpha_q}} C_x^{\frac{\alpha_x}{1-\alpha_q}}.$$

Proof. Postulate $V(q_j) = Aq_j$, so $\dot{V}(q_j) = 0$. From the first-order conditions,

$$\begin{aligned}
l_q &= \left(\frac{\alpha_q\chi A\lambda q^{\gamma_x+\gamma_q}\bar{q}^{\bar{\gamma}_x+\bar{\gamma}_q}}{w}\right)^{\frac{1}{1-\alpha_q}} l_x^{\frac{\alpha_x}{1-\alpha_q}} \\
&\Rightarrow \alpha_x l_x^{\alpha_x-1}\left((\chi A\lambda)^{\frac{1}{1-\alpha_q}}\left(\frac{\alpha_q\bar{q}}{w}\right)^{\frac{\alpha_q}{1-\alpha_q}} q^{\frac{\gamma_x+\gamma_q}{1-\alpha_q}}\bar{q}^{\frac{\bar{\gamma}_x+\bar{\gamma}_q-\alpha_q}{1-\alpha_q}} l_x^{\frac{\alpha_x\alpha_q}{1-\alpha_q}} + \chi_b q^{\gamma_b+\gamma_x}\bar{q}^{\bar{\gamma}_b+\bar{\gamma}_x}\right) = w
\end{aligned}$$

Suppose, $l_x = C_x\frac{q}{\bar{q}}$ and $l_q = C_q\frac{q}{\bar{q}}$ and replace on the previous equations to obtain a system of equations where constants C_x and C_q are implicitly defined,

$$\begin{aligned}
C_x^{1-\alpha_x} &= \frac{\alpha_x\bar{q}}{w}\left((\chi A\lambda)^{\frac{1}{1-\alpha_q}}\left(\frac{\alpha_q\bar{q}}{w}\right)^{\frac{\alpha_q}{1-\alpha_q}} C_x^{\frac{\alpha_x\alpha_q}{1-\alpha_q}} + \chi\chi_b\right) \\
C_q &= \left(\frac{\alpha_q\chi A\lambda\bar{q}}{w}\right)^{\frac{1}{1-\alpha_q}} C_x^{\frac{\alpha_x}{1-\alpha_q}}
\end{aligned}$$

Note the system has a unique solution, $\tilde{C}_{lx}^*(A)$ and $\tilde{C}_{lq}^*(A)$ given,

$$\frac{\alpha_x\alpha_q}{1-\alpha_q} < (1-\alpha_x) \iff 0 < 1-\alpha_x-\alpha_q.$$

$\square$

C The Implications of High-Speed: Entry

The reason for focusing on entry is twofold. First, entry represents a key indicator of activity in general (Cunningham et al., 2021) and in a given technology market, and thus allows researchers to evaluate market-level effects of firm-level activity. Second, and central to our analysis, is the role that speed plays in "protective" innovation. The relative increase value that goes to large firms is something not captured at the market level, in many cases, because large firms use quantity (or speed) instead of quality to create patent thickets and reduce downstream entry.

So far, we have focused on the decisions at the firm level and the allocation of inventors. Given our previous empirical evidence, we can furnish a quantitative model with counterfactuals on the speed and quality of firms' innovations, but what are the market-level effects of these different allocations?

We focus on the interaction between speed, quality, and entry in order to link these forces. In particular, as innovation moves from slower to faster firms, how will it affect the entry of other firms in the technology? We perform this analysis by looking at the class level.

We perform regressions that look at the change in entry at the class level

Table C.1: Speed vs Entrants

	(1) Log Number Entrants	(2) Log Number Entrants
Lag Top 1 Speed (Stock)	-0.557***	-0.091***
	(0.124)	(0.026)
Log Lag Number of patents	0.868***	0.701***
	(0.048)	(0.088)
Quarter FE	No	Yes
Class FE	No	Yes
Observations	14,801	14,801
R-squared	0.806	0.961

Note: The table presents a weighted OLS regression at class level of the lag number of entrants against the lag speed of top 1 firms (measured as stock), weighted by the lag number of patents. Controls include the log of the lag number of patents by industry class, and Column (2) includes additional quarter and industry class fixed effects. Standard errors clustered by industry class are reported in parentheses. ***, **, and * denote statistical significance at the 1, 5, and 10 percent level, respectively.

Table C.2: Quality vs Entrants

	(1) Log Number Entrants	(2) Log Number Entrants
Lag Top 1 Quality (Stock)	0.274***	0.078
	(0.079)	(0.064)
Log Lag Number of patents	0.791***	0.684***
	(0.047)	(0.089)
Quarter FE	No	Yes
Class FE	No	Yes
Observations	15,997	15,996
R-squared	0.784	0.960

Note: The table presents a weighted OLS regression at class level of the lag number of entrants against the lag quality of top 1 firms (measured as stock), weighted by the lag number of patents. Controls include the log of the lag number of patents by industry class, and Column (2) includes additional quarter and industry class fixed effects. Standard errors clustered by industry class are reported in parentheses. ***, **, and * denote statistical significance at the 1, 5, and 10 percent level, respectively.

Table C.3: Speed vs Entrants

	(1) Log Number Entrants	(2) Log Number Entrants
Lag Top 1 Speed (Stock)	-0.020**	-0.053***
	(0.009)	(0.019)
Log Lag Number of patents	0.603***	0.640***
	(0.090)	(0.060)
Pre-95	Yes	No
Post-95	No	Yes
Quarter FE	Yes	Yes
Class FE	Yes	Yes
Observations	6,515	8,285
R-squared	0.928	0.971

Note: The table presents a weighted OLS regression at class level of the lag number of entrants against the lag speed of top 1 firms (measured as stock), weighted by the lag number of patents. Controls include the log of the lag number of patents by industry class, and Column (2) includes additional quarter and industry class fixed effects. Standard errors clustered by industry class are reported in parentheses. ***, **, and * denote statistical significance at the 1, 5, and 10 percent level, respectively.

Table C.4: Quality vs Entrants

	(1) Log Number Entrants	(2) Log Number Entrants
Lag Top 1 Quality (Stock)	0.016	0.157**
	(0.028)	(0.076)
Log Lag Number of patents	0.588***	0.656***
	(0.088)	(0.056)
Pre-95	Yes	No
Post-95	No	Yes
Quarter FE	Yes	Yes
Class FE	Yes	Yes
Observations	7,444	8,552
R-squared	0.927	0.971

Note: The table presents a weighted OLS regression at class level of the lag number of entrants against the lag quality of top 1 firms (measured as stock), weighted by the lag number of patents. Controls include the log of the lag number of patents by industry class, and Column (2) includes additional quarter and industry class fixed effects. Standard errors clustered by industry class are reported in parentheses. ***, **, and * denote statistical significance at the 1, 5, and 10 percent level, respectively.