Technology-Driven Market Concentration through Idea Allocation*

Yueyuan Ma† Shaoshuang Yang‡

This Version: March 2026

Abstract

This paper documents that market concentration falls when frontier-shifting technologies surge. To explain this, we develop a quantitative-theoretic model where idea holders choose between launching startups and joining incumbents, yielding two predictions: (1) startup entry rises during technology-wave peaks, and (2) within incumbents, idea quality assortatively matches with firm size. Both predictions match the data and explain declining concentration during breakthroughs. Quantitatively, the breakthrough slowdown predicts faster average firm-quality growth, but this is more than offset by slower net entry. This slowdown largely accounts for market concentration's upward trend and fluctuations.

Keywords: technology waves, growth, HHI, startups, incumbent firms.

*We would like to thank Ufuk Akcigit, Thomas Chaney, Murat Celik, Emin Dinlersoz, Jesus Fernandez-Villaverde, April Franco, Ted Frech, Jeremy Greenwood, Jeremy Pearce, Gerard Hoberg, Ayse Imrohoroglu, Boyan Jovanovic, Aakash Kalyani, Yueran Ma, Swapnika Rachapalli, Peter Rupert, Mehmet Yorukoglu, and seminar participants at UBC, UCR, USC, UF, UPenn, the NBER PIE conference, the Tepper-Laef conference, the Duke Macro Jamboree, and the Cement Workshop for helpful comments. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2125 (CBDRB-FY24-P2125-R11032, CBDRB-FY26-P2125-R12802)

†Affiliation: University of California, Santa Barbara. Email: yueyuanma@ucsb.edu.

‡Affiliation: The Chinese University of Hong Kong, Shenzhen. Email: yangshaoshuang@cuhk.edu.cn.

1 Introduction

The rise in U.S. market concentration and the slowdown in growth since the early 2000s have raised significant concerns. The period has witnessed declining firm entry and an expansion of productive incumbent firms (Autor et al. (2020)). The existing literature attributes this shift to the increasing market power of incumbents, which strategically deter the entry of new businesses (Cunningham, Ederer and Ma (2021); Akcigit and Goldschlag (2023)), thereby hindering the diffusion of innovative ideas. This paper offers a new perspective, providing empirical evidence and structural analysis to show that technology waves play an important role in shaping growth and concentration by reallocating innovative ideas between incumbents and entrants. The slowdown in technological breakthroughs contributes to slower growth and rising concentration, consistent with the view that good ideas are becoming harder to find (Bloom et al. (2020)).

Using USPTO, Longitudinal Business Database (LBD), and Compustat data, this paper documents several empirical patterns suggesting that technological waves are linked to market concentration through the allocation of innovative ideas among firms.

1. The aggregate breakthrough intensity experiences a declining trend, punctuated by medium-run cyclical waves.

2. Market concentration, as measured by the Herfindahl-Hirschman Index (HHI) of firm employment, payroll, or sales, exhibits both a rising trend and a cyclical pattern that is negatively correlated with the waves of the aggregate breakthrough intensity.

3. The allocation of innovation between startups and incumbents tracks technology waves: at the peaks, a larger share of patents are forming in new businesses, while at the troughs, a larger share of patents come from incumbent firms.

4. Among patents from incumbent firms, there is a positive relationship between idea quality, measured by patent citations, and firm size.

5. Conditional on a patent's scientific value, incumbent firm size increases its economic value, reflecting incumbent-inventor synergies. However, the economic return to scientific value declines during technology waves, highlighting adoption frictions for breakthrough ideas within incumbents.

These patterns indicate that the technology waves affect firm entries and the way new ideas combine with firms of different sizes.

Motivating by the empirical patterns, this paper proposes a theory about inventors' choice of where to contribute the value of their ideas, and how it links the technology

waves, growth, and market concentration. The technology waves are assumed to be random aggregate shocks governing the idea quality distribution in a period. Each inventor is endowed with an idea of idiosyncratic quality drawn from this distribution. The inventor needs to choose between forming a new firm of a random size (à la Romer (1990)) with a partner or joining an incumbent firm and climbing the quality ladder (Aghion and Howitt (1990); Grossman and Helpman (1991)). In the case of the latter, she must also decide on the size of the incumbent firm to join. These choices jointly determine new firm entry, quality upgrading by incumbents, and ultimately the firm size distribution. Simulations of the calibrated model show that a slowdown in breakthrough innovation lowers aggregate growth, even as average firm quality grows faster. Moreover, the concentration dynamics implied by technology waves are strongly positively correlated with their empirical counterparts.

The model features three ingredients that determine whether an inventor launches a startup or joins an incumbent. First, while a startup's idea value depends solely on its quality, joining an incumbent requires integrating with a legacy technology base. Because new ideas are imperfect substitutes for existing technologies, incumbents face adoption frictions, in the spirit of Jovanovic and Rousseau (2014), especially for high-quality ideas. Second, incumbents offer commercialization synergies through superior scale, which raises the economic value of adopted ideas. Third, unobservable inventor effort creates a principal-agent problem requiring wage-plus-equity contracts. However, following Holmstrom (1989) and Aghion and Tirole (1994), equity provides weaker incentives in large firms because their aggregate payoffs are less exposed to innovation-related shocks. Ultimately, inventors trade off trade off startups' lower adoption frictions, lack of synergy, and stronger incentive alignment against incumbents' greater synergy but higher adoption frictions and, for large firms, weaker incentives. This multidimensional trade-off microfounds the equilibrium allocation of ideas.

The model delivers two main predictions. First, at the peaks of technology waves, a larger share of inventors choose to launch new firms to develop their ideas. Second, among inventors that choose to do R&D in incumbent firms, there is positive assortative matching between idea quality and firm size. Therefore, firms already with a larger size attract ideas of higher value. Both predictions align with the data and together imply lower market concentration near the peak of a technology wave. The surge in startup formation expands the mass of firms, and because startups are not subject to positive assortative matching, they counterbalance the tendency of large incumbents to accumulate ever more valuable ideas and further expand.

To quantify how technology waves affect economic growth and market concentration through the allocation of new ideas, we calibrate the model and run counterfactual

simulations that vary the idea-quality distribution while holding all other parameters fixed. The simulations are conducted year by year starting in 1986, the first peak of the technology waves in our sample. We then compute annual growth rates and decompose them into (i) growth driven by improvements in average firm quality and (ii) growth driven by net expansion in firm mass. We also generate time paths for two key moments: (1) the Herfindahl–Hirschman Index (HHI) of firm size, and (2) the ratio of ideas developed in new firms relative to incumbent firms. The simulated paths are compared to their empirical counterparts.

The evolution of technology waves leads to a lower aggregate growth rate after the 2000s, consistent with data. Although average firm quality grows faster, the growth rate in firm mass decreases more. The simulated paths of the HHI and idea allocation closely track their empirical counterparts. In particular, the model-generated HHI reproduces the observed upward trend. Moreover, the correlation between the detrended simulated and actual HHI ranges from 0.658 to 0.793 across alternative HHI measures. Taken together, these results indicate that technology waves are a quantitatively important driver of both growth and market concentration.

To decompose how idea allocation shapes market concentration, we distinguish between an extensive margin (firm mass) and an intensive margin (matching between idea quality and firm size). We trace the extensive margin through the allocation of ideas between startups and incumbents, with the detrended model-implied new-to-incumbent ratio correlating 0.763 with its empirical counterpart. We isolate the intensive margin with a counterfactual that holds firm mass fixed while changing only the matching pattern. The results show that startup formation mitigates positive assortative matching, lowering both the level and the upward trend of market concentration.

To capture rising barriers to entry, growing patent thickets, and the increasing dependence of new businesses on services provided by incumbents, we introduce a startup tax that rises over time and is rebated to incumbents in proportion to their revenue. Re-simulating the model shows a larger rise in market concentration, especially after the 2000s, with more ideas flowing toward incumbents.

Related Literature

In the endogenous growth literature, innovation is commonly modeled as either variety expansion (as in Romer (1990)) or quality-ladder improvement (as in Aghion and Howitt (1990); Grossman and Helpman (1991); Klette and Kortum (2004); Akcigit and Kerr (2018)). However, their distinct implications for growth and market concentration have not been systematically examined. On the growth side, this paper shows that inventors endogenously choose between these two innovation modes by trading off adoption

frictions against synergies. This trade-off shapes the aggregate growth rate and varies over technology waves. On the market concentration side, if each firm produces a single variety, variety expansion predicts an increase in the firm mass and a decrease in concentration. In contrast, quality-ladder growth has ambiguous implications for concentration, affected by how idea quality is matched with firm size. This paper provides both empirical evidence and a theoretical mechanism, driven by the trade-off between synergy and incentive alignment, that generates positive assortative matching. As a result, market concentration increases when ideas join incumbents.[1]

This paper explores the macroeconomic implications of groundbreaking technologies. Empirically, Bowen III, Frésard and Hoberg (2023) show that during eras of rapidly evolving technology, startups increasingly remain independent rather than being acquired. Dinlersoz, Dogan and Zolas (2024) document a massive surge in AI business applications after 2016. On the theoretical front, Greenwood and Yorukoglu (1997), Atkeson and Kehoe (2007), and Jovanovic and Rousseau (2014) have formalized how incumbent firms face frictions when adapting to frontier technologies and explored their aggregate growth impacts. Our quantitative-theoretic model extends this literature by linking these adoption frictions directly to the endogenous ownership and allocation of innovative ideas. By doing so, we establish that market concentration is a fundamental, endogenous outcome of technological revolutions. Ultimately, this paper demonstrates that alongside the high-frequency business cycles influenced by productivity fluctuations (Kydland and Prescott (1982)), the economy may also be governed by low-frequency cycles driven by technology waves.

Our empirical and theoretical analyses highlight the role of technology waves in shaping where inventors conduct R&D and, in turn, market concentration. This offers a new perspective on the connection between idea allocation and concentration. Prior studies (e.g., Cunningham, Ederer and Ma (2021); Akcigit and Goldschlag (2023)) emphasize that incumbents may strategically acquire and shelve external innovations to deter competition, contributing to a decline in patent quality. Our findings do not contradict these conclusions, but instead provide a complementary mechanism. Specifically, this paper implies that technology waves and market concentration can mutually reinforce each other, strengthening the negative association between the two.

This paper offers a new explanation for the rise in U.S. market concentration, especially since the early 2000s. While this trend has coincided with gains in allocative efficiency and productivity (Autor et al. (2020); Ganapati (2021)), it has also been

[1]Related work on the interaction between the firm size distribution and growth includes Chatterjee and Rossi-Hansberg (2012); Perla, Tonetti and Waugh (2021); Fons-Rosen, Roldan-Blanco and Schmitz (2021); Hopenhayn, Neira and Singhania (2022); Cassiman and Veugelers (2006); Bena and Li (2014); Akcigit, Celik and Greenwood (2016); Cavenaile, Celik and Tian (2019); Liu and Ma (2021); Ma (2026); Yang (2026).

accompanied by declining firm entry and a widening gap between large and small firms (Akcigit and Ates (2023); Olmstead-Rumsey (2019)), and growing innovation difficulty (Bloom et al. (2020)). We propose that a slowdown in radical technological breakthroughs helps reconcile these observations. Our Tech Wave Index reveals that major waves of innovation peaked in the mid-1980s and mid-1990s, with a prolonged 17-year lull until a resurgence in the early 2010s. During this stagnation, inventors increasingly turned to incumbent firms, enabling large incumbent firms to further expand.

The rest of the paper is organized as follows. Section 2 introduces measures of the technology waves, market concentration, and the allocation of ideas, and subsequently presents their patterns. Section 3 constructs a model where inventors make decisions between initiating new ventures or joining established incumbents at specific sizes. We derive predictions about the mapping between the quality of inventors' ideas and their optimal choices. Section 4 calibrates the model. Section 5 simulates the model to evaluate the degree to which technology waves can account for changes in growth rates and market concentration through the idea allocation channel. Section 6 concludes.

2 Empirical Patterns

This section exhibits empirical patterns of the technology waves, market concentration, and the choices of the inventors on where to contribute their ideas. A description of the data used in this section is provided in Appendix A.

2.1 Technology waves

Technology waves capture the aggregate breakthrough intensity of new technologies over time. At the peak of a technology wave, there is a surge of mass innovation driven by high-quality, novel, and influential ideas. At the trough, most developments are incremental improvements upon existing ones, while next-generation technologies may still be in an early stage. Importantly, because this dynamic captures the combined novelty and influence of all technologies in each year, it is distinct from the diffusion process of a technology from emergence to full economy-wide dispersion, which usually takes decades, as documented by Atkeson and Kehoe (2007) and Kalyani et al. (2025).

2.1.1 Measurement

We measure the aggregate intensity of technological breakthroughs using three methods. The first is the "RETech" index developed by Bowen III, Frésard and Hoberg (2023), an aggregate-level measure that leverages patent-text vocabulary to quantify whether

overall innovation is concentrated in rapidly evolving versus stable technological areas.[2] The second relies on the patent-importance metric developed by Kelly et al. (2021), which evaluates importance at the individual patent level by scaling a patent's forward similarity to future patents by its backward similarity to prior art. We calculate the aggregate breakthrough intensity by taking the average of this metric across all patents granted in each year. However, the measure is truncated at 2005. Therefore, we construct a third, citation-based Technology Wave Index (TWI), extending the conceptual logic of Kelly et al. (2021) beyond 2005. TWI captures the same ratio of forward-looking impact to backward-looking reliance using patent citations. Specifically, for each year t, we define:

$$\text{TWI}_t = \frac{\sum_{i \in I_t} \sum_{s=0}^{5} \text{Forward Citations}_{i,t+s}}{\sum_{i \in I_t} \sum_{s=0}^{5} \text{Forward Citations}_{i,t+s} + \sum_{i \in I_t} \sum_{s=0}^{5} \text{Backward Citations}_{i,t-s}}, \tag{1}$$

where I_t is the set of new patents granted in year t. The numerator sums the forward citations (citations received from others) that each new patent accumulates over the subsequent five years. The denominator scales this by the sum of these forward citations plus the backward citations (citations made to prior patents) granted within the previous five years. The fixed five-year windows ensure that cohorts are compared on common ground, mitigating right-censoring issues. To smooth short-term volatility, we apply a three-year moving average:

$$\text{TWI_avg}_t = \frac{1}{3} \sum_{h=-1}^{1} \text{TWI}_{t+h}.^{3} \tag{2}$$

Panel A of Figure 1 plots our citation-based TWI_avg_t alongside the aggregate text-based series from Kelly et al. (2021) and Bowen III, Frésard and Hoberg (2023). They exhibit remarkably similar trends and fluctuations with largely synchronized peaks and troughs. This co-movement confirms that the identified technology waves reflect deep structural shifts in the innovation frontier rather than measurement artifacts. Across all measures, major technological breakthroughs occur in the mid-1980s, the mid-1990s, and the early 2010s, while the periods around 1990 and the mid-2000s represent phases of predominantly incremental innovation.[4]

One concern of patent-based measures is whether they are driven by mechanical factors, such as raw patent volume. To address this concern, Panel B of Figure 1 reports

[2] A related measure by Kalyani (2024) uses new patent phrases to capture creativity.

[3] The smoothing technique does not alter the underlying cyclical patterns, as shown in Appendix B.1.

[4] Although our primary analysis focuses on the US economy, we compute the citation-based TWI for six high-patenting European countries using PATSTAT. As shown in Figure 15 in Appendix B.5, the broad secular decline in breakthrough intensity from the 1980s through the 2010s is also present across these major European patenting economies.

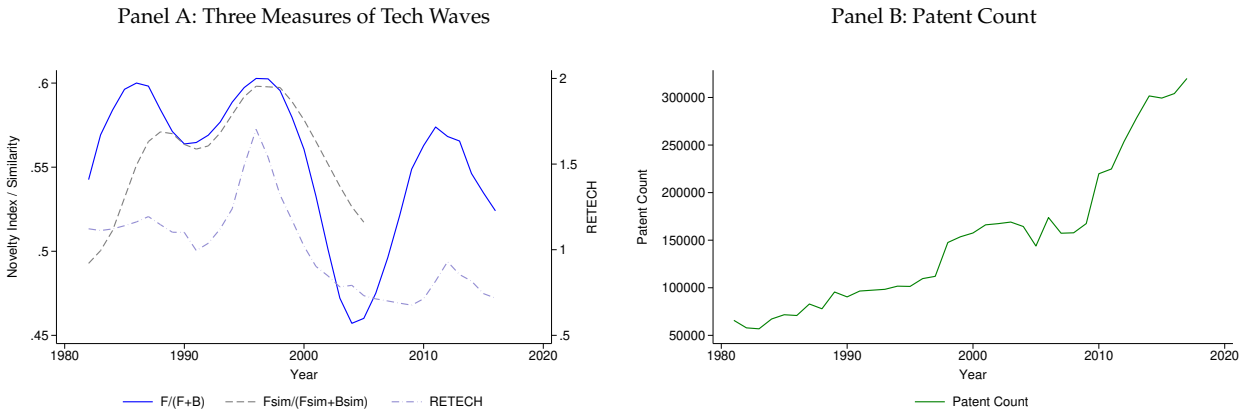


Figure 1: Technology Waves

Notes: Panel A illustrates three measures of technology waves. The blue solid curve is the citation-based measure defined in Equation 1, while the gray dashed curve is based on the text-similarity metric in Kelly et al. (2021); both use the left y-axis. The purple dash-dot curve shows the "RETech" index from Bowen III, Frésard and Hoberg (2023), which assesses novelty via the prevalence of growing vocabularies in patent descriptions, plotted on the right y-axis. Panel B shows the total granted patent count by year.
Sources: USPTO patent and citation data.

the total number of granted patents over time, demonstrating that the cyclical fluctuations in these wave indices are driven by distinct surges in relative patent quality and novelty. Field-level indices, constructed by the first digit of the International Patent Classification (IPC), are reported in Figure 13 in Appendix B.2 and exhibit both broad macroeconomic co-movement and expected heterogeneity across specific fields.

2.1.2 Contributors to the Tech Waves

Which classes of technology contributed to the three wave peaks? Who were the major applicants for breakthrough patents—incumbents, startups, or public institutions?

We address these questions in two steps (with full details in Appendix B.3 and B.4). First, we decompose TWI_t into contributions by three-digit IPC classes. The results highlight shifting technological frontiers: *Medical or Veterinary Science and Hygiene* is the largest contributor to the first peak (1985–1987), while *Computing; Calculating or Counting* becomes the dominant driver in both the second (1995–1997) and third (2010–2012) peaks, alongside *Electric Communication Technique*.

Second, to characterize who generates the most prominent breakthroughs, we cross-reference the "historically significant patents" compiled by Kelly et al. (2021). Breaking down applicants by institution type reveals that breakthrough patents are filed by a remarkably broad set of organizations: 41% by incumbents, 26% by startups, and 20% by universities. Because the creation of breakthroughs is widely distributed, our structural

framework will focus on how ideas are allocated after breakthroughs rather than on who creates them.

2.2 Market Concentration

Our primary measure of market concentration is the Herfindahl-Hirschman Index (HHI), constructed using two datasets: the Census Bureau's Longitudinal Business Database (LBD) for all U.S. employers, and Compustat Fundamentals Annual for publicly listed U.S.-headquartered industrial firms. We compute annual industry-level HHIs by summing squared firm market shares—defined by employment or payroll at the 6-digit NAICS level in the LBD, and by sales at the 4-digit SIC level in Compustat. We then aggregate these into an economy-wide HHI using industry employment or sales weights, respectively. Finally, we apply a three-year moving average to smooth the series.[5]

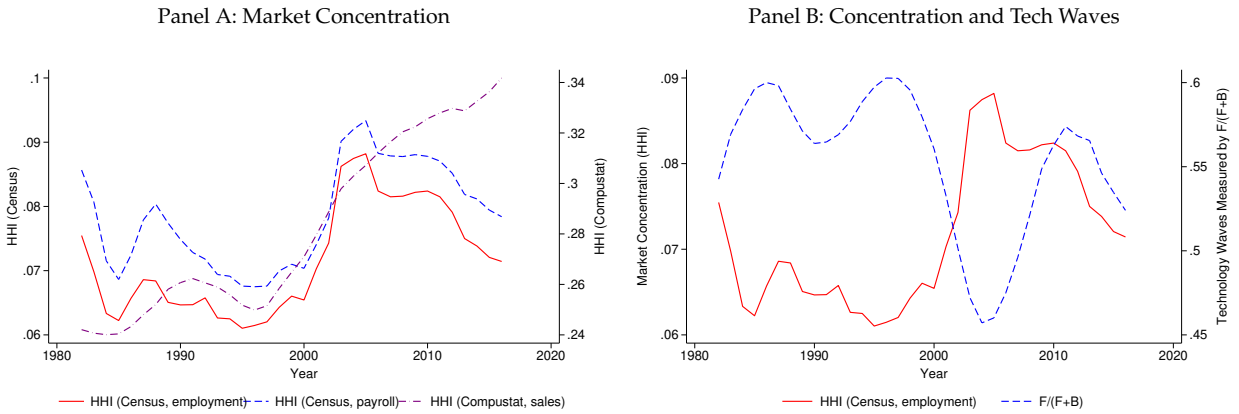


Figure 2: Technology Waves and Market Concentration

Notes: Panel A displays the annual HHI for employment (red solid curve), payroll (blue dash curve), and sales (purple dash-dot curve). The first two are derived from the LBD dataset from the Census, while the last one is based on Compustat. Panel B illustrates the technology waves alongside the trend of market concentration over time. The red solid curve shows the HHI for employment. The blue dashed curve, following the methodology defined in this paper, represents the relative ratio of forward citations to the sum of forward and backward citations. The scale of y-axes are shown on the left and right, respectively.
Sources: Longitudinal Business Database (LBD), Compustat Fundamental Annuals, and USPTO patent data.

Panel A of Figure 2 displays the annual Herfindahl-Hirschman Index (HHI) for firm employment, payroll, and sales, all of which exhibit similar trends and fluctuations, except for the last ten years. The pairwise correlations among these measures are high: 0.96 between employment and payroll, 0.73 between employment and sales, and 0.66 between payroll and sales. Panel B illustrates the relation between market concentration, measured by the employment-based HHI, and the technology waves measured by TWI.

[5]Figure 11 in Appendix B.1 plots the unsmoothed LBD employment HHI and Tech Wave Index, which exhibit a similar correlation.

The two series are negatively correlated. The technology waves exhibit a downward linear trend, while the HHI shows an upward trend. The cross-correlation between the detrended HHI (x_t) and the detrended technology waves (y_{t+k}) are -0.736, -0.775, and -0.720 respectively when k equals to 0, -1, -2. This suggests that changes in market concentration closely follow technology waves with a lag.[6]

To assess the robustness of market concentration patterns, the share of sales by top firms is calculated using the cleaned data series from Kwon, Ma and Zimmermann (2023), which is based on IRS data covering the entire population of U.S. corporations. Figure 16 in Appendix B.8 shows that the top sales shares exhibit similar upward trends and cyclical patterns to the HHI.[7]

A more granular analysis is conducted using regression analysis, as specified in Equation (3):

$$\text{HHI}_{st} = \beta_0 \text{TWI}_{st} + \beta_1 \text{Size}_{st} + \theta_s + \mu_t + \epsilon_{st}. \tag{3}$$

Here HHI_{st} is the Herfindahl–Hirschman Index for 6-digit NAICS industry s in year t based on the Census data. The right–hand side includes the industry's Tech Wave Index, industry fixed effects θ_s, and year fixed effects μ_t. To rule out the effect of technology waves on HHI through changes in market size (Campbell and Hopenhayn (2005)), we control for industry size, measured by total employment or payroll in each 6-digit NAICS industry. The Tech Wave Index is constructed at the 4-digit IPC level and mapped to 6-digit NAICS using a concordance that we build by tabulating, in the Census data, the industry composition of patent owners within each IPC category.

Instruments Because the Tech Wave Index may be correlated with unobserved shocks to concentration and reverse causality may exist, we construct a shift-share instrument that combines the U.S. innovation network with lagged technology waves. Using U.S. patent data, we construct a time-varying diffusion matrix $\Omega_{c' \to c,t,\tau}$ at the IPC-class level that maps citation flows from an upstream (cited) class c' in year $t - \tau$ to a downstream (citing) class c in year t, for lags $\tau \in \{1, \ldots, 5\}$. For each (c, t, τ) we expand patent citations to IPC$\times$IPC pairs with fractional weights, drop within–class links, aggregate to (c', c, t, τ), and normalize so that $\sum_{c' \neq c} \Omega_{c' \to c,t,\tau} = 1$. We then combine this diffusion structure with the Tech Wave Index at the IPC level, $TWI_{c',t-\tau}$, to form a shift–share instrument in the

[6] Cross-correlations for $k \in \{-3, -2, -1, 0, 1, 2, 3\}$ are reported in Table 11 in Appendix B.6. The strongest correlation, in absolute value, occurs at $k = -1$.

[7] The aggregate negative correlation between technology waves and market concentration is also evident in Europe, as shown by the declining trend of the Tech Wave Index in Appendix B.5 and the increasing market concentration across European countries, as measured by the Herfindahl-Hirschman Index (HHI) and top sales share in recent studies (e.g., Bighelli et al. (2023) and Ma, Zhang and Zimmermann (2024)).

spirit of Acemoglu, Akcigit and Kerr (2016) and Liu et al. (2025):

$$\lambda_{c,t} \;=\; \frac{1}{5}\sum_{\tau=1}^{5}\sum_{c'\neq c}\Omega_{c'\to c,t,\tau}\,TWI_{c',\,t-\tau}.$$

Intuitively, $\lambda_{c,t}$ loads each downstream class on lagged breakthrough intensity in technologically connected upstream classes. The instrument is mapped from IPC to NAICS using the same concordance. Following the exogenous shocks framework of Borusyak, Hull and Jaravel (2022), our key identifying assumption is that the exact timing and magnitude of aggregate technological breakthroughs in the *upstream* classes (c') are quasi-random with respect to downstream (c), sector-specific unobservables. Because downstream market participants cannot dictate the fundamental scientific progress of their upstream suppliers, this network-based instrument isolates the effect of technological shocks and effectively mitigates the problem of reverse causality.

The regression results are reported in Table 1. Columns (1)–(4) examine the effect of contemporaneous technological waves, columns (5)–(8) use a one-year lag for the waves, and columns (9)–(12) use a two-year lag. For each set of specifications, we report the OLS estimates first, followed by the 2SLS estimates, both with and without industry and year fixed effects. The corresponding first-stage results and F-statistic of the instrument are reported in Table 12 in Appendix B.7. Across all columns, higher breakthrough intensity predicts lower market concentration.

Table 1: Relationship between HHI and Tech Wave Index at the 6-digit NAICS code

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	OLS		2SLS		OLS		2SLS		OLS		2SLS	
TWI	-0.142***	-0.0284	-0.241***	-0.306**								
	(0.0343)	(0.0367)	(0.0850)	(0.139)								
TWI (1-year lag)					-0.137***	-0.0218	-0.238***	-0.367**				
					(0.0336)	(0.0363)	(0.0831)	(0.148)				
TWI (2-year lag)									-0.130***	-0.0555	-0.330**	-0.407***
									(0.0327)	(0.0341)	(0.144)	(0.147)
Observations	37500	37500	37500	37500	37500	37500	37500	37500	36500	36500	36500	36500
R-squared	0.011	0.69			0.011	0.69			0.01	0.701		
Industry size	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Industry FE	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes

Notes: Standard errors are clustered at the 6-digit NAICS and year level. Columns (1)–(4) examine the effect of contemporaneous technological novelty, columns (5)–(8) use a one-year lag, and columns (9)–(12) use a two-year lag. All specifications control for industry size. The inclusion of year and industry fixed effects is indicated in the last two rows. To comply with Census Bureau disclosure requirements, the number of observations is rounded to the nearest hundred. *** Significant at the 1 percent level; ** significant at the 5 percent level; * significant at the 10 percent level.

2.3 Allocation of Ideas

One potential link between the technology waves and the market concentration is idea holders' choices of where to do innovation. They can start their own businesses or contribute their innovation efforts to incumbent firms of selected sizes. This section describes the flow of the new ideas using the Census and USPTO patent data.

2.3.1 Entrants or Incumbent Firms

Though affiliations at idea inception are not directly available, it can be inferred by observing the firm age at the time of patent issuance, from the combined USPTO and LBD data. Specifically, when a patent is granted to a firm aged zero to five years, it indicates that the initial idea was developed in a startup five years earlier. In contrast, when a patent is granted to a firm older than five years, it implies that the idea was developed internally by an incumbent firm. The five-year cutoff reflects average USPTO application-to-grant lags (about 2–3 years as documentented by the USPTO) and a comparable research completion period. Under these assumptions, the ratio of ideas in new firms to ideas absorbed by incumbents is measured as:

$$\text{New-to-Incumbent Ratio}_t = \frac{\sum_{i \in I_{t+5}} \text{Granted in Firm(Age} \leq 5)_{i,t+5}}{\sum_{i \in I_{t+5}} \text{Granted in Firm(Age>5)}_{i,t+5}}, \tag{4}$$

where I_{t+5} denotes the set of patents granted five years after time t, and the indicators "Granted in Firm(Age$\leq$5)" and "Granted in Firm(Age$>$5)" equal one if patent i is granted to a firm aged five years or less, or older than five years, respectively.

As in previous analyses, a three-year moving average is applied to smooth the time series, and the resulting trend is shown in Figure 3.[8] For comparison, the technology waves measured by TWI is also plotted. A declining trend is observed in both series over time. In addition, the New-to-Incumbent Ratio exhibits pronounced cyclicality, with its peaks and troughs closely aligned with those of the technology waves. The cross-correlation between the detrended New-to-Incumbent Ratio (x_t) and the detrended Tech Wave Index (y_{t+k}) is computed across different time lags and reported in Table 11 in Appendix B.6. The maximum absolute correlation, 0.612, occurs at $k = 0$, indicating that the two series move in tandem.

[8] Figure 12 in Appendix B.1 presents the unsmoothed series for both the New-to-Incumbent Ratio and the Tech Wave Index, which display a similar correlation structure to their smoothed counterparts.

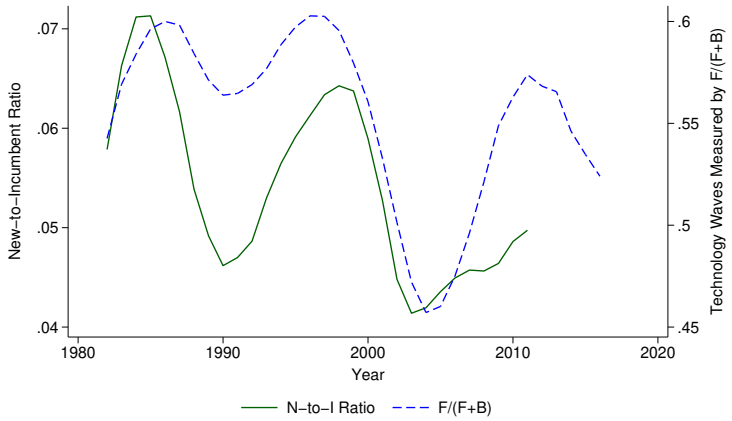


Figure 3: Technology Waves and Idea Allocation

Notes: This figure shows the technology waves and the idea allocation between new and incumbent firms over time. The blue dashed curve, based on the methodology defined in this paper, calculates the relative ratio of forward citations to the sum of forward and backward citations. The green solid curve displays the "New-to-Incumbent Ratio" defined in the paper, capture where new ideas contribute their value. The two curves have different y-axises, which are shown respective on the left and right.
Sources: Longitudinal Business Database (LBD) and USPTO patent and citation data.

2.3.2 Granular Relationship between Tech Waves and the New-to-Incumbent Ratio

The positive correlation between technology waves and the New-to-Incumbent Ratio is evident across most major technological fields, including human necessities, performing operations, chemistry, textiles, fixed constructions, physics, and electricity. Graphs by field illustrating this relationship are presented in Figure 17 in Appendix B.9. Analysis at a more granular level is conducted using a regression, as specified in Equation (5):

$$\text{New firm}_{ict} = \beta_0 \text{TWI}_{ct} + \beta_1 \text{Citation}_{ict} + \theta_c + \mu_t + \epsilon_{ict}. \tag{5}$$

The dependent variable is an indicator equal to one if a patent is granted to a new firm (age $\leq$ 5). We regress it on the Tech Wave Index measured at the 4-digit IPC (c) by year (t) level, controlling for the patent's scientific value proxied by the number of forward citations it receives within five years of issuance, and including IPC and year fixed effects. To address the endogeneity issue, we apply the shift-share instrument that combines the U.S. innovation network with lagged technology waves.[9] The regression results are reported in Table 2. Columns (1)–(2) present OLS estimates, while columns (3)–(4) report the 2SLS results. The first-stage regression results are presented in Table 12 in Appendix B.7. We report specifications both with and without IPC and year fixed

[9]Since in this regression, the Tech Wave Index is at the 4-digit IPC level, there is no need to map the Index to the 6-digit NAICS codes.

effects. Across all specifications, a significant positive relationship is found, indicating that idea holders are more likely to develop innovations in new firms when breakthrough technologies occur. This association remains robust after accounting for business-cycle fluctuations at the IPC and year levels.

Table 2: Starting New Businesses and Tech Wave Index at the Patent Level

	(1)	(2)	(3)	(4)
	OLS		2SLS	
TWI	0.0560***	0.014	0.127***	0.179***
	(0.00996)	(0.0141)	(0.0477)	(0.0514)
Ln(1+Citations)	0.0104***	0.0113***	0.00966***	0.0108***
	(0.000947)	(0.000926)	(0.000935)	(0.000952)
Observations	4032000	4032000	4032000	4032000
R-squared	0.003	0.013		
IPC FE	No	Yes	No	Yes
Year FE	No	Yes	No	Yes

Notes: Standard errors are clustered at the 4-digit IPC and year level. Columns (1) and (3) include no fixed effects. Columns (2) and (4) incorporate both IPC and year fixed effects. To comply with Census Bureau disclosure requirements, the number of observations is rounded to the nearest thousand. *** Significant at the 1 percent level; ** Significant at the 5 percent level; * Significant at the 10 percent level.

2.3.3 Size of Incumbent Firms

When inventors contribute ideas to incumbent firms, they also choose firm size. This link is examined using patents granted to firms at least five years old (the incumbent firms, as defined in Section 2.3.1). Idea quality is proxied by the number of forward citations each patent receives within five years of issuance, and firm size is measured by employment. We pool patents issued to both new and incumbent firms across years and compute citation quartiles to group patents by relative citation counts. Within each quartile, average assignee firm size (LBD employment) is computed, the lowest-quartile mean is normalized to one, and the resulting relative sizes are plotted in Figure 4.

The figure reveals a pattern of positive assortative matching between idea quality and firm size: higher-quality ideas are more likely to be matched with larger firms. However, we do not find this assortative matching pattern when calculating the average assignee size across citation quartiles for patents granted to firms younger than five years.

2.4 Economic Value of Patents over Tech Waves

Results so far suggest that incumbent firms are less attractive during the peak of tech waves. To explore the underlying channels, we perform patent-level regressions. We

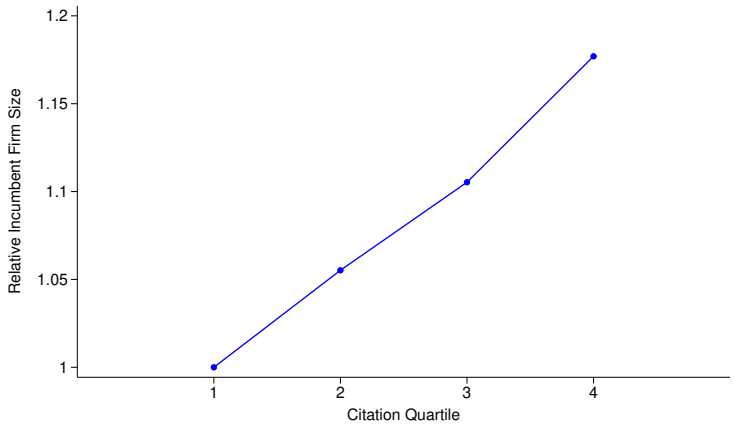


Figure 4: Mapping between Patent Citations and Incumbent Firm Size

Notes: This figure shows the mapping between inventors' idea quality and firm size if inventors opt to develop their ideas in incumbent firms. The idea quality is measured by the number of patent citations and is classified into four quartiles. The firm size is measured by the number of employees. The average employment of firms corresponding to the first citation quartile is normalized to be one.
Sources: Longitudinal Business Database (LBD) and USPTO patent and citation data.

use an extended version of data by Kogan et al. (2017). It leverages the stock market's response to patent news to estimate the private economic value of patents. Since the sample encompasses all patents granted to publicly listed firms in the US, it provides valuable insights into factors affecting the economic value of patents in incumbent firms over technology waves. The following regression analysis is conducted,

$$\begin{aligned}\ln(\text{Economic value}_{it}) =& b\ln(\text{Firm size}_{j(i)t}) + \iota\ln(1+\text{Citations}_{it}) + \\ & \phi\text{TWI}_t \times \ln(1+\text{Citations}_{it}) + \theta_{c(i)t} + \gamma_{j(i)t} + \epsilon_{it}.\end{aligned} \tag{6}$$

where i, j, c, and t are indexes for patents, firms, patent technology classes at the 4-digit IPC level, and years. The dependent variable corresponds to the economic value of the patents. Firm size is measured by sales of the firm to which the patent belongs. The number of forward citations received within five years of issuance is used to measure the scientific value of patents, serving as a proxy for idea quality. The interaction term between the Tech Wave Index and the citations captures the impact of technology waves on the relationship between the scientific and economic value of patents. The model controls for IPC-by-year fixed effects in all specifications and firm-by-year fixed effects in some specifications.

Table 3 presents the results using firm employment as the measure of firm size. Columns (1)-(3) exclude the technology wave measure, focusing solely on the properties of patents and firms. Column (4) displays results of Equation (6). In Columns (5), the

yearly Tech Wave Index is replaced by the IPC-by-year Tech Wave Index.

Table 3: Factors of Patent's Economic Value for Incumbent Firms

	Ln(Patent's Economic Value)				
	(1)	(2)	(3)	(4)	(5)
Ln(1+Employment)	0.329***				
	(0.064)				
Ln(1+Sales)		0.335***			
		(0.031)			
Ln(1+Citations)	0.085***	0.092***	0.004***	0.032**	0.013**
	(0.017)	(0.017)	(0.001)	(0.012)	(0.006)
Ln(1+Citations)$\times$TWI$_t$				-0.053**	
				(0.022)	
Ln(1+Citations)$\times$TWI$_{ct}$					-0.018*
					(0.010)
Observations	1,107,289	1,113,710	1,103,218	1,103,218	1,103,218
R-squared	0.369	0.470	0.884	0.884	0.884
IPC $\times$ Year FE	Yes	Yes	Yes	Yes	Yes
Firm $\times$ Year FE	No	No	Yes	Yes	Yes

Notes: Standard errors are clustered at the 4-digit IPC-by-year level. Columns (1)-(3) exclude the technological wave measure and focus solely on the property of the patents and firms. Columns (4) shows coefficients of the regression equation (6). Columns (5) replaces the yearly Tech Wave Index by the IPC-by-year Tech Wave Index. The regressions control for IPC-by-year across all specifications. The firm-by-year fixed effects are controlled in columns (3)-(5). *** Significant at the 1 percent level; ** Significant at the 5 percent level; * Significant at the 10 percent level.

Firm size has a significantly positive effect on the economic value of patents, given the idea quality. This suggests that the synergy between inventors and firms increases with firm size. Additionally, idea quality positively impacts the economic value of patents, but this impact diminishes with higher aggregate breakthrough intensity, as indicated by the negative coefficients of the interaction terms. This finding implies that, during periods of more front-shifting technologies, ideas, especially high-quality ones, realize lower value within incumbents.

3 Model

To formalize how technology waves drive market concentration, we develop a general equilibrium model with two groups of individuals—households and inventors—and two types of firms: intermediate goods producers and final goods producers. In each period, inventors draw ideas of heterogeneous quality. They decide either to establish new intermediate-goods firms (Romer-type innovation) or to join incumbent firms of chosen size (quality-ladder innovation). The economy is subject to aggregate shocks that vary the scale of the idea quality distribution. The shocks shift the relative payoff to inventors

across innovation types. Inventors' choices depend jointly on the aggregate shock and the quality of their individual ideas.

3.1 Preferences

There is a long-lived representative household in the economy. She works in the production sector, supplies one unit of labor to firms inelastically, and consumes final goods. She also owns all the firms in the economy. The household's utility function is

$$U_H = \int_0^\infty e^{-\rho t} log(C_H(t))dt, \tag{7}$$

where $\rho > 0$ is the discount rate and $C_H(t)$ is the consumption of the household.

Inventors are the ones who work in the R&D sector. In each period, there is a continuum of inventors of measure one. An inventor, with a short-lived lifespan of dt time periods, dedicates effort $e_I(t)$ to create innovations within either an incumbent firm or a new business. Simultaneously, they engage in consumption. Inventors are risk-averse and have a mean-variance utility flow:

$$U_I\left(c_I(t), e_I(t)\right) = \mathbb{E}\left(c_I(t)\right) - A\frac{var\left(c_I(t)\right)}{V(t|Q(t))} - R\left(e_I(t)\right)V(t|Q(t)), \tag{8}$$

where $c_I(t)$ is the consumption, $e_I(t)$ is the effort level, and $R\left(e_I\right)V(t|Q(t))$ is the associated cost. $Q(t)$ (defined below) is the aggregate quality in the economy at time t, and $V(t|Q(t))$ is the value of a firm whose size equals to the aggregate quality. The variance and cost are normalized by $V(t|Q(t))$ to keep the problem stable over time. Denote the inventors' aggregate consumption by $C_I(t)$, i.e., $C_I(t) = \int_0^1 c_{Ii}(t)di$.

3.2 Technology

The economy features two types of firms: intermediate goods producers and final goods producers,as in Akcigit and Kerr (2018). Both types of firms are owned by the household. The former hires inventors to create innovations, and produce intermediate goods. The latter assembles intermediate goods and produces final goods.

The final good producers produce final goods using a continuum of intermediate goods $j \in [0, N_F(t)]$:

$$Y(t) = \frac{1}{1-\beta}\int_0^{N_F(t)} q_j^\beta(t) y_j^{1-\beta}(t)dj. \tag{9}$$

In this function, $q_j(t)$ is the quality of the intermediate good j, and $y_j(t)$ is its quantity. $N_F(t)$ is the total mass of intermediate goods producers. We normalize the price of the

final good to be one in every period. The final good producers are perfectly competitive, taking the input prices as given.

The intermediate goods producers are a continuum of risk-neutral firms. Each firm produces one type of good, with a linear technology using only labor:

$$y_j(t) = Q(t) l_j(t), \tag{10}$$

where $l_j(t)$ is the labor input; $Q(t) = \int_0^{N_F(t)} q_j(t) dj$ is the aggregate quality level of the economy. It implies that improvement in both q_j (intensive margin) and N_F (extensive margin) has positive externalities (Romer, 1986; Aghion and Howitt, 1990). The cost is linear in wage $w(t)$, which intermediate firms take as given. The labor market satisfies the constraint:

$$\int_0^{N_F(t)} l_j(t) dj \leq 1. \tag{11}$$

3.3 Innovation

Innovation takes two forms: creating new intermediate-good varieties (*Romer-type*) or improving the quality of existing lines (*Quality-ladder*). *Romer-type* innovation is led by an inventor pairing with an outside, risk-neutral partner. *Quality-ladder* innovation occurs within incumbent firms, led either by an inventor partnering with the incumbent or by the incumbent itself after acquiring the idea. Because inventor effort is unobservable, a principal-agent problem featuring moral hazard arises whenever the inventor leads the project (Holmstrom, 1989; Aghion and Tirole, 1994). We assume an exogenous share $1 - h$ of ideas is best implemented entirely by incumbents; these are sold to incumbents of random size and follow a standard quality-ladder process. For the remaining share h, inventors choose to lead innovations in either startups or incumbents of an optimal size based on equilibrium contracts. This organizational choice trades off adoption frictions, incumbent synergies, and incentive alignment. Table 4 summarizes the features of each innovation method, which are detailed in the following sections.

Table 4: Features of Different Innovation Methods

Innovation Type	**Leader**	**Partner**	**Adoption Frictions**	**Synergy**	**Incentive Problems**	**Share**
Romer	Inventor	Partner	No	No	Yes	h
Quality-ladder	Inventor	Incumbent	Yes	Yes	Yes	
Quality-ladder	Incumbent	None	Yes	Yes	No	$1 - h$

3.3.1 Inventor-led R&D

The key to R&D management is finding the optimal incentive structure to motivate unobservable effort (Aghion and Tirole, 1994; Merges, 1999). Consistent with empirical evidence demonstrating that compensation structures directly drive both inventor effort and the allocation of innovation (Lerner and Wulf, 2007; Giarratana, Mariani and Weller, 2018), we explicitly model the principal-agent problem between inventors and firms. We assume all partners employ a widely adopted contract structure that combines a fixed wage for risk-sharing with an equity component for incentive alignment (Onishi, 2013).[10] Formally, inventors endogenously choose an effort level, e_I , which governs the innovation arrival rate according to a Poisson process:

$$\lambda(e_I) = \lambda_0 e_I. \tag{12}$$

Exerting effort is costly for inventors. The instantaneous cost of effort e_I is given by $R(e_I)V(t|Q(t))$, where $R(e_I) = \frac{1}{1+\delta} e_I^{\delta+1}$, an increasing and convex function of effort.[11]

Each inventor is endowed with a single innovative idea of quality z_0, drawn randomly from a distribution, $\Psi(z_0)$. If the inventor leads innovation in a startup, there is full pass-through from idea quality to realized innovation quality, $zQ(t)$, which is drawn from a uniform distribution,

$$\text{Uniform}\left((1-\phi)z_0Q(t),\ (1+\phi)z_0Q(t)\right). \tag{13}$$

On average, higher-quality ideas lead to higher-quality innovations, but the parameter $\phi \in (0,1)$ introduces randomness into this mapping by capturing variability in how idea quality translates into realized outcome.

If the inventor with idea quality z_0 leads co-innovation with an incumbent firm of quality q, the quality improvement, $x(z_0, \tilde{q}(t))Q(t)$, reflects characteristics of both the inventor and the firm. We model it as a stochastic draw from a uniform distribution:

$$\text{Uniform}\left((1-\phi)x_0(z_0,\tilde{q}(t))Q(t),\ (1+\phi)x_0(z_0,\tilde{q}(t))Q(t)\right), \tag{14}$$

where $\tilde{q}(t)$ denotes the incumbent's relative quality, defined as $\tilde{q}(t) = \frac{q}{Q(t)}$. For convenience in the subsequent analysis, we adopt the normalization $\tilde{m} = \frac{m}{Q(t)}$ for variable

[10] Because the inventor's effort level, e_I, is unobservable and unverifiable, contracts cannot be directly contingent on it. Consequently, firm equity is used to align the inventor's returns with the firm's.

[11] The innovation production function and cost specifications follow the growth-theory literature (Romer, 1990; Klette and Kortum, 2004; Akcigit and Kerr, 2018). In our calibration, we set $\delta = 1$, as is standard. We model inventors' R&D costs as a dis-utility. Equivalently, they can be represented as a resource cost that converts one-to-one into final goods. This modeling choice has no material effect on our results.

m when necessary. The effective idea quality, $x_0(z_0, \tilde{q}(t))$, depends positively on both the inventor's original idea quality and the relative firm size, i.e., $\frac{\partial x_0(z_0,\tilde{q}(t))}{\partial z_0} > 0$ and $\frac{\partial x_0(z_0,\tilde{q}(t))}{\partial \tilde{q}(t)} > 0$. The function $x_0(z_0, \tilde{q}(t))$ takes the form:

$$x_0(z_0, \tilde{q}(t)) = \chi(\tilde{q}(t))\gamma(z_0) = \left(\frac{\tilde{q}(t)}{q_0}\right)^b \left(B^\eta + z_0^\eta\right)^{\frac{1}{\eta}}, \ \eta < 1, \ b > 0.$$

The first term, $\left(\frac{\tilde{q}(t)}{q_0}\right)^b$, captures incumbents' synergy benefits, which rise with the firm's relative quality. The second term, $\left(B^\eta + z_0^\eta\right)^{\frac{1}{\eta}}$ with $\eta < 1$, is a constant-elasticity-of-substitution (CES) aggregator that captures the interaction between the inventor's idea and the incumbent's existing technology base B in the co-invention process. The CES structure with imperfect substitution between the new idea z_0 and the technological base, B, implies that the new idea cannot fully substitute for existing product lines within incumbent firms. So, idea quality does not fully pass through to realized economic value, reflecting frictions in adopting new technologies. This reduced-form approach to modeling compatibility costs follows the spirit of Jovanovic and Rousseau (2014).

Joining Incumbent Firms: Incumbent intermediate firms compete à la Bertrand to attract inventors by offering compensation packages consisting of wage rate $T_I V(t|Q(t))$ and equity a. The firm's quality, and equivalently its size, evolves stochastically:

$$dq = -qdJ_{\text{exit}} + x(z_0, \tilde{q}(t))Q(t)dJ_{R\&D}(e_I), \tag{15}$$

where J_{exit} and $dJ_{R\&D}(e_I)$ are two jump processes with Poisson arrival rates at τ and $\lambda(e_I) = \lambda_0 e_I$, respectively. The first jump process represents the exogenous exit, while the second one is related to the in-house innovation. The magnitude of the jump is $x(z_0, \tilde{q}(t))Q(t)$, the stochastic value of an innovation produced by inventor z_0 in a firm with a normalized size $\tilde{q}(t)$ when the aggregate technology level is $Q(t)$.[12]

The inventor's compensation $c_I(a, T_I, q, z_0, Q(t))$ is:

$$c_I(a, T_I, q, z_0, Q(t))dt = T_I V(t|Q(t))dt + adV(j, t|q, z_0), \tag{16}$$

where $dV(j, t|q, z_0)$ is the firm value change from t to $t + dt$ when hiring an inventor of quality z_0. Following a principal-agent framework, the compensation package includes, in addition to a fixed wage, an opportunity for the inventor to purchase a fraction a of

[12]The parameter b in $x_0(z_0, \tilde{q}(t))$, estimated from the data, governs how strongly the economic value depends on firm quality, $q(t)$, relative to the aggregate quality, $Q(t)$.

equity at the fair price. At the end of each period of length dt, the inventor sells this equity back to the firm to finance consumption. In doing so, the inventor bears part of the risk from exit shocks and innovation shocks. A firm j with quality q solves the following problem when an inventor with idea quality z_0 leads the innovation process:

$$\begin{aligned}\Omega_I(z_0,q,t) = \max_{a,T_I} &(1-a)\,\lambda_0 e_I\left[\int V(j,t|q+x(z_0,\tilde{q}(t))Q(t))dx - V(j,t|q)\right] - T_I V(t|Q(t)) \\ \text{s.t. } e_I &= \arg\max\left\{U_I\left(c_I(a,T_I,q,z_0,Q(t)),e_I\right)\right\} \\ &U_I\left(c_I(a,T_I,q,z_0,Q(t)),e_I\right) \geq \bar{u}\left(z_0,t\right) \\ &(1-a)\,\lambda_0 e_I\left[\int V(j,t|q+x(z_0,\tilde{q}(t))Q(t))dx - V(j,t|q)\right] - T_I V(t|Q(t)) \geq 0\end{aligned} \tag{17}$$

where $\Omega_I(z_0,q,t)$ denotes the expected payoff. The instant change rate in a firm's expected payoff from hiring the inventor equals the increase in firm value from the innovation accruing to the original shareholders (all shareholders other than the inventor), net of the wage and equity granted to the inventor. Here, $V(j,t|q)$ denotes the firm's value without R&D.

The first constraint is the inventor's incentive-compatibility condition: given the contract a, T_I, she chooses effort e_I to maximize expected utility $U_I\left(c_I(a,T_I,q,z_0,Q(t)),e_I\right)$. The second is the inventor's participation constraint, requiring utility at least $\bar{u}(z_0,t)$, where the outside option is endogenously pinned down by Bertrand competition. The third is the firm's participation constraint, ensuring the firm is not worse off from hiring the inventor.

Inventor's utility function can be expressed as:

$$\begin{aligned}U_I\left(c_I(a,T_I,q,z_0,Q(t)),e_I\right) =& \mathbb{E}\left(c_I(a,T_I,q,z_0,Q(t))\right) - A\frac{\text{Var}\left(c_I(a,T_I,q,z_0,Q(t))\right)}{V(t|Q(t))} \\ &- R\left(e_I\right)V(t|Q(t)).\end{aligned} \tag{18}$$

Starting up a New Business: If an inventor starts a new business with an outside partner, the partner solves an optimization problem analogous to that of an incumbent firm (Equation (17)). The key differences are that there is neither synergy nor adoption friction, so $q = 0$ and $x_0 = z_0$. We assume the partner's outside option value is zero, so its participation constraint also requires nonnegative net payoff. The inventor then chooses her effort level to maximize utility.[13]

[13] A detailed exposition of the startup problem is provided in Appendix C.1.

3.3.2 Firm-led R&D

If an incumbent firm leads the innovation, it purchases an idea from an inventor and selects its own effort level e_F. The innovation's success is then governed by an instantaneous Poisson arrival rate:

$$\lambda(e_F) = \lambda_0 e_F. \tag{19}$$

Exerting effort is also costly for firms. The instantaneous cost of effort e_F is given by $R(e_F)N_F(t)V(t|q(t))$, where $R(e_F) = \frac{1}{1+\delta}e_F^{\delta+1}$.[14] Both the cost function $R(\cdot)$ and the arrival rate $\lambda(\cdot)$ take the same functional forms as in the inventor-led innovation case. The total R&D cost borne by firms is:

$$R_F(t) = \int_0^{N_F(t)} R(e_F(t, q(j))N_F(t)V(t|q(t))dj. \tag{20}$$

Since firms bear both the costs and the benefits of innovation, they are fully incentivized to choose the effort level that maximizes their payoff. If the purchased idea has quality z_0, a successful innovation raises the firm's quality from $q(t)$ to $q(t) + x_F N_F(t)q(t)$, where $x_F N_F(t)q(t)$ is drawn from a uniform distribution similar to (13):

$$\text{Uniform}\left((1-\phi)\kappa x_{F0}N_F(t)q(t),\ (1+\phi)x_{F0}N_F(t)q(t)\right). \tag{21}$$

where $x_{F0} = \gamma(z_0) = \left(B^\eta + z_0^\eta\right)^{\frac{1}{\eta}}$. $\kappa > 0$ captures the scale effects of a firm-led innovation. As compensation, the firm makes a fixed transfer $T_F N_F(t)V(t|q(t))$ to the inventor. The firm maximizes the expected value of innovation by solving the following problem,

$$\begin{aligned}\Omega_F(z_0, q, t) = \max_{e_F, T_F} \lambda_0 e_F &\left[\int V(j, t|q + \kappa x_{F0}N_F(t)q)dz - V(j, t|q)\right] \\ &- \frac{1}{1+\delta}e_F^{\delta+1}N_F(t)V(t|q(t)) - T_F(t)N_F(t)V(t|q(t)).\end{aligned} \tag{22}$$

This problem is similar to a standard quality-ladder model. As in the inventor-led case, we assume that the transfer T_F equals the firm's expected innovation value.[15]

$$\lambda_0 e_F[\int V(j, t|q + \kappa x_{F0}N_F(t)q)dz - V(j, t|q)] - \frac{1}{1+\delta}e_F^{\delta+1}N_F(t)V(t|q(t)) = T_F N_F(t)V(t|q(t)) \tag{23}$$

[14] The introduction of $N_F(t)V(t|q(t))$ makes the cost and benefit of firm-led R&D comparable with innovator-led R&D.

[15] An alternative assumption is that the incumbent captures all surplus. The resulting expressions are of a similar functional form, since firm value remains linear in size.

3.4 Timeline

At birth, each inventor observes the initial quality z_0 of her idea. A share $1-h$ of ideas are sold to an incumbent of random size. For the remaining share h, an outside partner—who also observes z_0—offers a contract to jointly establish a new intermediate firm. At the same time, incumbent firms observe their corresponding $x_0\left(z_0, \tilde{q}\right)$ and extend employment contracts to the inventor. After evaluating all offers, the inventor either joins her preferred incumbent with relative quality $\tilde{q}^*(z_0, t)$ or launches a startup with the outside partner.

The inventor's decision rule is:

$$u\left(z_0, t\right) = \max\{U_I\left(c_I(z_0, q^*(z_0, t), t), e_I(z_0, q^*(z_0, t), t\right), U_I\left(c_I(z_0, 0, t), e_I(z_0, 0, t)\right)\}. \quad (24)$$

3.5 Entry and Exit

Incumbent intermediate firms face an exogenous exit rate τ, which is independent of their size and is a risk unrelated to innovation.[16] After a firm exits the market, the technology it owns is distributed to new entrants randomly. Entry of new intermediate firms occurs through successful innovations by inventors launching startups with outside partners. The entry rate is denoted by $\lambda_I(t)$:

$$\lambda_I(t) = \int_{z_0 \in \{q^*(t)=0\}} h\lambda_0 e_I\left(z_0, 0, t\right) d\Psi(z_0). \quad (25)$$

Upon entry, an entrant draws a relative quality level $\tilde{q}_I$ from the distribution $f_{I,t}\left(\tilde{q}_I\right)$, where it satisfies

$$f_{I,t}(\tilde{q}_I) = \frac{\lambda_I(t)}{\tau N_F(t)} f_t(\tilde{q}), \tilde{q}_I = \frac{\tau N_F(t)}{\lambda_I(t)}\tilde{q}.$$

where $f_t\left(\tilde{q}\right)$ is the incumbents' quality distribution. Therefore, the total quality flow of exiting incumbents is equal to the total quality flow of entrants,

$$\lambda_I(t)\int q_I f_{I,t}(q_I) dq_I = \lambda_I(t)\int \tilde{q}_I Q(t) f_{I,t}(\tilde{q}_I) d\tilde{q}_I = N_F(t)\int \tau\tilde{q} Q(t) f_t(\tilde{q}) d\tilde{q} = \tau Q(t). \quad (26)$$

The new firm incurs a cost equal to the value of the technology associated with its drawn quality, $V(t|q_I)$. It then augments a quality update by applying the innovation of the inventor.

The law of motion for the mass of intermediate-goods firms is

$$\dot{N}_F(t) = \lambda_I(t) - \tau N_F(t). \quad (27)$$

[16] A modest negative correlation between exit rates and firm size does not affect the model's predictions.

3.6 Household's and Firms' Choices

The household maximization problem reveals the relationship between the growth rate, $g(t)$, and the time discount factor,

$$\frac{\dot{C}_H(t)}{C_H(t)} = r(t) - \rho. \tag{28}$$

The final good producer chooses $\{y_j(t)\}_j$ to maximize its profit using the technology described in Section 3.2, which yields the demand function faced by intermediate goods producers: $p_j(t) = q_j(t)^{\beta} y_j(t)^{-\beta}$. The intermediate good producers engage in monopolistic competition.[17] Their FOC yields,

$$y_j(t) = q_j(t)\left(\frac{Q(t)\,(1-\beta)}{w(t)}\right)^{\frac{1}{\beta}}, l_j(t) = \frac{y_j(t)}{Q(t)}, p_j(t) = \frac{w(t)}{Q(t)\,(1-\beta)}. \tag{29}$$

In each period, the labor market clearing condition pins down the wage:

$$w(t) = (1-\beta)\,Q(t). \tag{30}$$

Thus, both the production output $y_j(t)$ and profit $\pi_j(t)$ are linear in quality,

$$y_j(t) = q_j(t), \pi_j(t) = \beta q_j(t), Y(t) = Q(t). \tag{31}$$

The HJB equation of an incumbent firm's value is

$$\begin{aligned} r(t)V(j,t|q) - \dot{V}(j,t|q) &= \pi_j(t) + \Omega_I(z_0,q,t) + \Omega_F(z_0,q,t) - \tau V(j,t|q) \\ &= \pi_j(t) - \tau V(j,t|q) \end{aligned} \tag{32}$$

The second line follows from the scarcity of inventors and the assumption that competition for their ideas ensures that inventors capture the full value of innovations. Since $\pi_j(t) = \beta q_j(t)$, the value of an intermediate firm is linear in its quality, i.e.,

$$V(j,t|q) = \nu(t) q. \tag{33}$$

Solving the firm-led innovation problem yields the incumbent firm's effort, $e_F = (\lambda_0 \kappa x_{F0})^{\frac{1}{\delta}}$, which is invariant to time and firm size[18]. The arrival rate, $\lambda(e_F) = \lambda_0^{\frac{1+\delta}{\delta}} (\kappa x_{F0})^{\frac{1}{\delta}}$, is constant across firms. The transfer to the inventor, $T_F(t) N_F V(t|q(t)) =$

[17] The profit maximization problem and the solution process of the final good and intermediate good producers are shown in Appendix C.2.

[18] See Appendix C.3 for details.

$\frac{\delta}{1+\delta}\left(\lambda_0 \kappa x_{F0}\right)^{\frac{1+\delta}{\delta}} \nu(t) N_F q(t)$, is linear in $\nu(t)$. Inventors' ideas are sold to firms randomly.

Solving the inventor-led innovation problem shows that the inventor's effort, e_I, is independent of time when (33) holds. The inventor's consumption, $c_I(a, T_I, q, z_0, Q(t))$, and utility, $U_I(c_I, e_I)$, are both linear in $\nu(t)$. The inventor's optimal firm choice is invariant to time and depends only on parameters.[19]

3.7 Aggregate Variables

The final goods are consumed by the household and inventors, or invested into R&D by firms. The resource constraint of the economy is:

$$Y(t) = C_H(t) + C_I(t) + R_F(t), \tag{34}$$

Inventors' total consumption, $C_I(t)$, can be expressed as the following equation by applying the law of large numbers :

$$\begin{aligned} C_I(t) = \nu(t) Q(t) & \int_{z_0 \in \{z_0 | \tilde{q}^* > 0\}} \left(h \lambda_0 e_I\left(z_0, \tilde{q}^*\right) x_0\left(z_0, \tilde{q}^*\right) + (1-h) T_F\right) d\Psi(z_0) \\ + \nu(t) Q(t) & \int_{z_0 \in \{z_0 | \tilde{q}^* = 0\}} \left(h \lambda_0 e_I\left(z_0, \tilde{q}^* = 0\right) z_0 + (1-h) T_F\right) d\Psi(z_0). \end{aligned} \tag{35}$$

Firms' R&D costs, $R_F(t)$, can be expressed by

$$R_F(t) = \frac{1-h}{1+\delta} \nu(t) Q(t) \int_{z_0} e_F^{\delta+1} dz_0. \tag{36}$$

The growth is from a single source—innovation. By the law of large numbers, aggregate growth can be written as:

$$\begin{aligned} g(t) = \frac{\dot{Q}(t)}{Q(t)} = & \int_{z_0 \in \{z_0 | \tilde{q}^* > 0\}} \left(h \lambda_0 e_I\left(z_0, \tilde{q}^*\right) x_0\left(z_0, \tilde{q}^*\right) + (1-h) \lambda_0 e_F \kappa x_{F0}\right) d\Psi\left(z_0\right) \\ & + \int_{z_0 \in \{z_0 | \tilde{q}^* = 0\}} \left(h \lambda_0 e_I\left(z_0, \tilde{q}^* = 0\right) z_0 + (1-h) \lambda_0 e_F \kappa x_{F0}\right) d\Psi\left(z_0\right). \end{aligned} \tag{37}$$

The aggregate growth rate only depends on the allocation of inventor-led innovations across firms, which is invariant to time, implying that:

$$g(t) = g.$$

[19] See Appendix C.4 for a formal proof.

Note that the growth rate of the aggregate quality can be decomposed into two parts:

$$\frac{\dot{Q}(t)}{Q(t)} = \frac{\dot{N}_F(t)}{N_F(t)} + \frac{Q(t)\dot{/}N_F(t)}{Q(t)/N_F(t)} \tag{38}$$

The first term captures growth from the extensive margin (firms' mass), and the second term captures growth from the intensive margin (average quality).

3.8 Equilibrium

This section characterizes an equilibrium of the economy in which aggregate variables (Y, C, R, w, Q) grow at a rate g.

Definition An equilibrium of this economy is the mapping between $\tilde{q}^*(t)$ and z_0, the allocation $\left(\left\{y_j^*(t)\right\}_j, Y^*(t), C_I^*(t), C_H^*(t)\right)$, the prices $\left(w^*(t), r^*(t), \left\{p_j^*(t)\right\}_j\right)$, the growth rate g^*, the entry rate $\lambda_I^*(t)$, and the measure of firms $N_F^*(t)$, such that (1) for any $j \in [0,1]$, $y_j^*(t)$ and $p_j^*(t)$ satisfy Equation (29); (2) the wage $w^*(t)$ satisfies Equation (30); (3) the interest rate $r^*(t)$ satisfies Equation (28); (4) the measure of the intermediate producers $N_F^*(t)$ satisfies Equation (27); (5) the mapping between $\tilde{q}^*(t)$ and z_0 comes from inventors' optimal choices; (6) the entry rates $\lambda_I^*(t)$ satisfy Equation (25); (7) R&D spending $C_I^*(t)$ and $R_F^*(t)$ satisfy Equation (35) and Equation (36), respectively; (8) the aggregate output Y^* satisfies Equation (31); (9) the aggregate consumption C_H^* satisfies Equation (34); and (10) the growth rate g satisfies Equation (37).

Proposition 1. *Consider the above-described economy starting with an initial condition $Q(0) > 0$. There exists a unique equilibrium. In this equilibrium, growth is always balanced, i.e., $g(t) = g^*$, $\forall t$. The aggregate quality, $Q(t)$, aggregate output, $Y(t)$, and aggregate consumptions, $C_H(t)$ and $C_I(t)$, all grow at the rate g^*.*

Lemma 1. *There is no transition dynamics.*

The proofs of Proposition 1 and Lemma 1 are provided in Appendix C.5. In what follows, we omit the time subscript t whenever doing so does not create confusion.

3.9 Inventor's Choice in a Closed Form

To illustrate the intuition and main channels in the model, this section presents a simplified setup that characterizes the inventor's trade-off in closed form. One modification is introduced in this setup: the realized innovation quality, $xQ(t)$, is drawn from a distribution with mean $x_0\left(z_0, \tilde{q}\right) Q(t)$ and second moment $e_I^{-1} x_0\left(z_0, \tilde{q}\right)^2 Q(t)^2$,

instead of from a uniform distribution $U\left(\left(1-\phi\right)x_0\left(z_0,\tilde{q}\right)\nu,\left(1+\phi\right)x_0\left(z_0,\tilde{q}\right)\nu\right)$. For startups, x and $x_0\left(z_0,\tilde{q}\right)$ are replaced by z and z_0, respectively. This modification captures an environment in which inventor effort reduces innovation uncertainty. The variance of the inventor's consumption becomes

$$\text{Var}\left(c_I\left(a,T_I,\tilde{q},z_0\right)\right)=a^2\nu^2Q(t)^2\left(\underbrace{\tau\tilde{q}^2}_{\text{Var}\tilde{V}(\tilde{q})}+\underbrace{\lambda_0\left(x_0\left(z_0,\tilde{q}\right)^2\right)}_{\text{Var(innovation)}}\right).$$

Innovation-related uncertainty no longer depends on the effort level.

The parameter δ in the effort cost function, $R(e_I)$, is assumed to take the value of 1, consistent with empirical estimates in the literature (e.g., Akcigit and Kerr (2018)). Using backward induction, firms anticipate that the inventor will chooses the following effort:

$$e_I=\lambda_0 a x_0\left(z_0,\tilde{q}\right).$$

A higher inventor equity share, a, or a higher expected innovation quality, $x_0!\left(z_0,\tilde{q}\right)$, induces greater inventor effort. In both cases, the marginal return to effort rises, making additional investment more attractive.[20]

The firm's problem in Equation 17 yields,

$$a^*=\frac{1}{1+2\frac{A}{\lambda_0}\left(\frac{\tau\tilde{q}^2}{\lambda_0 x_0\left(z_0,\tilde{q}\right)^2}+1\right)} \tag{39}$$

When $b<1$, the optimal equity share a decreases with firm size $\tilde{q}$. This reflects two opposing forces that jointly determine a: the commercialization value $x_0\left(z_0,\tilde{q}\right)$ and the non-innovation shock, captured by the exit risk $\tau\tilde{q}^2$. Firm size affects these components in opposite directions. On the one hand, larger firms provide greater synergy, which raises $x_0\left(z_0,\tilde{q}\right)$ and makes it more valuable to incentivize inventor effort, pushing a up. On the other hand, larger firms face higher exit risk, so offering less equity shields the inventor from uncertainty unrelated to innovation, pushing a down. The net effect depends on which channel dominates. When $b<1$, as in the calibration, the exit-risk channel dominates, and larger firms optimally offer a smaller equity share to the inventor.

Denote the optimal compensation scheme by (a,T_I). Under Bertrand competition, the wage component T_I^* is pinned down by the firm's zero-profit condition,

$$T_I^*=\left(1-a^*\right)\lambda_0 e_I x_0\left(z_0,\tilde{q}\right).$$

[20]The derivation of the optimal contract is provided in Appendix C.6.

Upon reviewing all contracts, an inventor with idea quality z_0 chooses the firm $\tilde{q}$ she would like to work for. If joining an incumbent, the optimal firm size is

$$\tilde{q}^* = \left(\frac{\left(2A\lambda_0 + \lambda_0^2\right)\left(\gamma\left(z_0\right)\right)^2 b}{2A\tau q_0^{2b}\left(1-2b\right)} \right)^{\frac{1}{2-2b}}.$$

Proposition 2. *When $b < 0.5$, there is positive assortative matching when ideas join incumbents: $\frac{\partial \tilde{q}^*}{\partial z_0} > 0$.*[21]

When b is low, synergy does not grow without bound as firm size increases, and the countervailing force of a lower equity share guarantees existence and uniqueness of the solution. The model predicts positive assortative matching among incumbent firms: higher-quality innovations are more likely to be developed in larger firms. This arises for two reasons. First, better ideas gain more from commercialization synergy. Second, they generate relatively more innovation-related uncertainty, which makes inventors less vulnerable to incentive problems.

Proposition 3. *If z_0 is drawn from a distribution $\Psi_{new}(z_0)$ that first-order stochastically dominates $\Psi(z_0)$, so ideas are better on average, a larger share of inventors launch startups.*[22]

Higher idea quality raises innovation output in both incumbents and startups. However, incumbents must integrate new ideas into an existing technology base, which lowers their realization potential, especially for high-quality ideas. As a result, when ideas are more frontier-shifting on average, as at the peak of a technology wave, starting up new firms becomes relatively more attractive.

Proposition 4. *When $b < \frac{\min\left(z_0^{-\eta}\right)}{\min\left(z_0^{-\eta}\right)+\max\left(B^{-\eta}\right)}$, there exists a cutoff $\bar{z}_0\left(B\right)$, such that all inventors with $z_0 < \bar{z}_0\left(B\right)$ join incumbent firms.*[23]

Both incumbents' adoption frictions and synergies increase with idea quality. If synergy gains do not rise too sharply with firm size, i.e., $b < \frac{\min!\left(z_0^{-\eta}\right)}{\min!\left(z_0^{-\eta}\right)+\max!\left(B^{-\eta}\right)}$, adoption frictions dominate, and inventors with high-quality ideas choose to start new businesses.

4 Calibration

We calibrate the full model to target the average US economy from 1982 to 2016. Patents are used as a surrogate for innovations. An innovation's idea quality, denoted by z_0,

[21] See Appendix C.7 for the proof.
[22] See Appendix C.8 for the proof.
[23] See Appendix C.9 for the proof.

and the realized value, $x\nu Q(t)$ ($z\nu Q(t)$ in the context of a startup), correspond to the patent's citation (scientific importance) and the patent's economic value, respectively. Additionally, we assume that the idea quality z_0 follows the Pareto distribution characterized by a scale factor z_m and a shape factor α.

4.1 Identification

Parameters in the model are categorized into two groups. The first group is calibrated by a prior information from the aggregate statistics or the literature. The second group is calibrated by estimation from the micro-level data or through the model. Table 5 reports the parameters in the first group, $(\rho, \beta, \tau, A, \delta)$. The discount rate, ρ, is set to 0.02 to match the average interest rate in the sample period. The production function quality share, β, is 0.109, following Akcigit and Kerr (2018). The firm exit rate, τ, is 0.09, targeting the average exit rate of firms during our sample period based on the Business Dynamics Statistics (BDS).[24] The scale parameter of the innovation arrival rate, λ_0, is normalized to 1 [25]. The risk aversion parameter, A, and the effort cost elasticity, δ, are set to be 0.5 and 1, respectively, which are commonly used in the literature (Hall and Van Reenen, 2000).

Table 5: Parameter Values from a Priori Information

Parameter	Description	Value	Identification
ρ	Discount rate	0.02	Interest Rate
β	Production function quality share	0.109	Firm profitability
τ	Exo. exit rate	0.09	BDS
λ_0	Innovation arrival rate	1	Normalization
A	Risk aversion	0.5	Risk aversion
δ	Effort cost elasticity	1	Effort cost elasticity

Notes: This table shows parameter values from the literature or direct estimation.

We jointly calibrate the ten remaining parameters in the second group, $(\alpha, z_m, \phi, \eta, B, b, q_0, h, \kappa)$, using the minimum distance method inspired by Lentz and Mortensen (2008). The parameters, along with their corresponding moments are shown in Table 6.

Growth Rate—Innovation is the sole driver of growth in the model. The scale parameter governing firm-led innovation, κ, is central to firms' innovation intensity and thus important to aggregate growth. A higher κ increases the step size of firm-led innovation, raising the overall growth rate. We calibrate κ so that the model-implied

[24] The BDS data is compiled from the Longitudinal Business Database (LBD) by the Census Bureau.

[25] In the model, λ_0 enters only multiplicatively, either as $\lambda_0 x_0(z_0, \tilde{q})$ or as $\lambda_0 \mathbb{E}\left[x_0(z_0, \tilde{q})^2\right]$, which implies that λ_0 is not separately identified from $x_0(z_0, \tilde{q})$.

Table 6: Parameters from the Minimum Distance Estimation

Para.	Description	Identification
α	Shape of idea quality distribution	SD/Mean of patent citations
z_m	Scale of idea quality distribution	Aggregate economic value of innovations
ϕ	Innovation value dispersion	SD/Mean of economic value cond. on cites
B	Maturity of technology	Technology Tech Wave Index
η	Elasticity of substitution	Estimation of the economic-value function
b	Exponent of the synergy function	Estimation of the economic-value function
q_0	Scale of the synergy function	New-to-incumbent ratio
h	Share of inventor-led innovation	Firm size ratio by cite quartiles (Q4/Q1)
κ	Scale of the firm-led innovation	Entrant/incumbent citation ratio

Notes: Parameters in this table are jointly calibrated to minimize the distance between the model and data moments.

aggregate growth rate equals 1.7%, matching the average annual U.S. growth rate in GDP per capita from 1982 to 2016.[26]

The S.D.-to-Mean Ratio of Patent Citations—This ratio captures the dispersion in patent citations observed in the data, which reflects the underlying dispersion in inventors' idea quality. The parameter α is the primary driver of this dispersion. Specifically, the standard deviation-to-mean ratio of the idea distribution is given by $\frac{1}{\sqrt{\alpha(\alpha-2)}}$. Although the patents recorded in the USPTO data represent only successful innovations—a selected subset of all ideas—the dispersion in patent citations remains heavily influenced by α. We construct the citation distribution by pooling all patents granted since 1976 and their corresponding citations recorded by the USPTO, and compute the standard deviation-to-mean ratio, which is approximately 2.784.

Aggregate Economic Value of Innovation—The economic value of innovations is modeled as a direct contribution to firm value. In the model, this value follows a uniform distribution with a mean determined by the underlying idea quality (scientific value), z_0. Given α, the average idea quality is governed by the scale parameter of the Pareto distribution, z_m. Accordingly, z_m is calibrated to match the ratio of average economic value of patents to average firm value.[27] Following the methodology of Kogan et al.

[26] We compute annual growth rates from the Federal Reserve Economic Data (FRED) as the trend component obtained by applying the HP filter to GDP per capita.

[27] Mathematically, the model expresses the ratio as a function of the underlying idea-quality distribution,

$$\begin{aligned}\frac{\mathbb{E}(\text{patent value})}{\mathbb{E}(\text{firm value})} &= \frac{\int_{z_0\in\{z_0|\tilde{q}^*>0\}} h\lambda_0 e_I\left(z_0,\tilde{q}^*\right) x_0\left(z_0,\tilde{q}^*\right)\nu Q}{\nu Q/N_F} + \frac{\int_{z_0\in\{z_0|\tilde{q}^*=0\}}(1-h)\,\lambda_0 e_F \kappa x_{F0}\nu N_F \int q dF(q)\, d\Psi\left(z_0\right)}{\nu Q/N_F} \\ &= N_F\int_{z_0\in\{z_0|\tilde{q}^*>0\}} h\lambda_0 e_I\left(z_0,\tilde{q}^*\right) x_0\left(z_0,\tilde{q}^*\right) + N_F\int_{z_0\in\{z_0|\tilde{q}^*=0\}}(1-h)\,\lambda_0 e_F\kappa x_{F0}\, d\Psi\left(z_0\right).\end{aligned} \tag{40}$$

(2017), patent economic value is inferred from stock market reactions using an extended dataset (1926–2022) linking U.S. patents to CRSP and Compustat data. While the raw aggregate patent value is 0.0069 times the average firm value for public firms, applying the statistical model by Yang (2026) to account for the broader firm population adjusts this ratio to 0.027. To balance these estimates, we adopt a calibrated value of 0.008.

S.D.-to-Mean Ratio of Innovations' Economic Value Conditional on Citations—Conditional on idea quality, the economic value of innovation is uniformly distributed, implying a standard-deviation-to-mean ratio of $\frac{\phi}{\sqrt{3}}$. Using the same sample as in the calibration of z_m, we estimate the standard-deviation-to-mean ratio of patent economic value while controlling for patent citations. In the data, this ratio is approximately 0.416.

Estimation of the Function of Patents' Economic Values—Both the elasticity of substitution between the technology base and new ideas, η, and the regression coefficient of innovations' economic value on incumbent firm size, b, are estimated directly from the data using nonlinear least squares. For each patent i granted to firm j in year t, we observe its economic value $x_{ijt}\nu Q(t)$, the firm's quality q_{jt}, the average technology base (proxied by average backward citations per patent) B_t, and the patent's forward citations z_{ijt}. The structural model for the underlying patent value is

$$x_{ijt}\nu Q(t) = \alpha\left(\frac{q_{jt}}{Q(t)}\right)^b \left(B_t^\eta + z_{ijt}^\eta\right)^{\frac{1}{\eta}} \nu Q(t)\varepsilon_{ijt},$$

where b and η are parameters of interest and ε_{ijt} is assumed to be a multiplicative disturbance with $E[\ln \varepsilon_{it} \mid q_{jt}, B_t, z_{ijt}] = 0$.

Taking logarithms and absorbing constants into an intercept and year fixed effects we obtain the estimating equation

$$\ln(x_{ijt}\nu Q(t)) = c + b \ln q_{jt} + \frac{1}{\eta} \ln\left(B_t^\eta + z_{ijt}^\eta\right) + \delta_t + u_{ijt},$$

where c collects α and ν, δ_t is a year fixed effect capturing $Q(t)$, and u_{ijt} is an error term satisfying $E[u_{ijt} \mid q_{ijt}, B_t, z_{ijt}, t] = 0$. Let the full parameter vector collect all unknowns,

$$\theta = \left(c,\, b,\, \eta,\, \{\delta_t\}_{t\in\mathcal{T}}\right),$$

and define the residual

$$u_{ijt}(\theta) = \ln(x_{ijt}\nu Q(t)) - c - b \ln q_{jt} - \frac{1}{\eta} \ln\left(B_t^\eta + z_{ijt}^\eta\right) - \delta_t.$$

We estimate θ by nonlinear least squares, i.e.,

$$\hat{\theta} = \arg\min_{\theta} S(\theta), \qquad S(\theta) = \sum_{ijt} u_{ijt}(\theta)^2.$$

At the optimum $\hat{\theta}$, the first-order conditions with respect to every component of θ are satisfied:

$$\frac{\partial S(\theta)}{\partial \theta_k} = 2 \sum_i u_i(\theta) \frac{\partial u_i(\theta)}{\partial \theta_k} = 0, \qquad k = 1, \ldots, \dim(\theta).$$[28]

The elasticity η and the scale factor b, included in θ, are identified in the estimation.

Technological Wave Index—Technology waves are defined as the ratio of forward citations to the sum of forward and backward citations for all patents granted in a given year. In the model, the existing technology base, B, which enters the realization potential of new ideas, $\gamma(z_0) = (z_0^\eta + B^\eta)^{1/\eta}$, corresponds to the stock of backward citations, representing the maturity of the existing technologies. The value of B is calibrated so that the model-generated average forward-to-backward citation ratio, $\int z_0 \, d\Psi(z_0) / (B + \int z_0 \, d\Psi(z_0))$, matches the average Tech Wave Index between 1982 and 2016, where $\int z_0 \, d\Psi(z_0)$ corresponds to the total forward citations of available ideas.

New-to-Incumbent Ratio—The scale parameter in the synergy function, $\tilde{q}_0$, governs the payoff from contributing an idea to an incumbent relative to starting a new venture, and thus shapes inventors' choice between incumbents and startups. We calibrate $\tilde{q}_0$ to match the "New-to-Incumbent Ratio" defined in Section 2.3.1, targeting its average value over 1982–2016.

Firm Size Ratio: Fourth-to-First Quartile of Patent Citations—The model predicts that, conditional on joining an incumbent, inventors with higher-quality ideas should match with larger firms. This positive sorting is weakened by firm-led innovation, since ideas are sold to incumbents of random size. As the share of firm-led innovation increases, the relationship between idea quality and firm size becomes flatter. To calibrate the share of inventor-led innovation, h, we compute in the model the average incumbent firm size by patent-citation quartile (restricting to patents developed within incumbents) and match the implied fourth-to-first quartile firm-size ratio to its empirical counterpart. [29]

[28] These conditions can be interpreted as a system of moment restrictions. Defining the moment vector

$$g(W_{ijt}, \theta) = u_{ijt}(\theta) \frac{\partial u_{ijt}(\theta)}{\partial \theta},$$

where $\frac{\partial u_{ijt}(\theta)}{\partial \theta}$ is the gradient of the residual with respect to all components of θ, the population moment conditions are

$$E\big[g(W_{ijt}, \theta_0)\big] = 0.$$

Thus, the nonlinear least squares estimator $\hat{\theta}$ can be viewed as a just-identified GMM estimator that jointly estimates $(c, b, \eta, \{\delta_t\}_{t \in \mathcal{T}})$ by enforcing the orthogonality conditions implied by the model.

[29] Within incumbents, innovations originating from an idea of quality z_0 (with $z_0 \in \mathcal{Z}$) fall into two

4.2 Calibration Results

Table 7 reports the model-generated moments and their counterparts in the data. Overall, the model matches the targeted moments closely. The resulting parameter values are reported in Table 8.

Table 7: Moments

Identification Moment	Data	Model
Growth rate	0.017	0.017
S.d.-to-mean ratio of patent citations	2.78	2.78
Aggregate innovation value	0.008	0.008
S.d.-to-mean ratio of innovation value cond. on citations	0.416	0.416
Elasticity of Substitution between past Knowledge Stock and New Ideas	-0.31	-0.31
Technology Wave Index	0.54	0.54
Regression coefficient of innovation value on firm size	0.32	0.32
New-to-incumbent ratio	0.054	0.054
Firm size ratio by fourth-to-first-quartile of citations	1.18	1.19

Notes: This table compares the moments generated from the calibrated model and the data. In general, the model-generated moments match the data well.

Table 8: Estimated Parameter Values

Parameter	Description	Value
α	Shape of idea quality distribution	2.10
z_m	Scale of idea quality distribution	0.086
ϕ	Dispersion of patent economic value	0.72
η	CES elasticity of substitution	-0.31
B	Technology base	0.13
b	Exponent of the synergy function	0.32
q_0	Denominator of the synergy function	9.2E-4
h	Share of inventor-led innovation	0.39
κ	Scale of the firm-led innovation	3.21

Notes: Parameters in this table are jointly calibrated to minimize the distance between the model and data moments.

categories: (i) inventor-led innovation and (ii) firm-led innovation. The average size of incumbent firms that own patents with z_0 in a set $\mathcal{Z}$ is

$$\mathbb{E}_{z_0\in\mathcal{Z}}(q) = \int_{z_0\in\mathcal{Z}} hq^*(z_0)\lambda_I^*(z_0)d\Psi(z_0) + \int_{z_0\in\mathcal{Z}} (1-h)\,\lambda_F^*(z_0)\int qf(q,t)dqd\Psi(z_0).$$

Here, $\int q, f(q,t), dq$ is the average incumbent firm size. It enters the firm-led term because firm-led ideas are allocated randomly across incumbents according to the cross-sectional firm-size distribution $f(q,t)$.

Our estimates suggest that, relative to startups, incumbents realize lower innovation value on average because of incomplete pass-through. In particular, the ratio, $\frac{\int \gamma(z_0)\, d\Psi(z_0)}{\int z_0\, d\Psi(z_0)} = 0.090$, is significantly below one. At the same time, synergy increases value. The scale parameter in the synergy function is $\tilde{q}_0 = 9.2 \times 10^{-5}$, and the elasticity is -0.31, implying that an average-sized incumbent can generate roughly 18 times the commercialization value of a startup through synergy alone. Once incomplete pass-through is taken into account, this net advantage falls to about 1.61 times.

Technology waves shape idea allocation along both the extensive and intensive margins. The extensive margin determines how many inventors found new businesses, while the intensive margin governs the firm size an inventor selects when joining an incumbent. Figure 5 illustrates both margins by plotting the optimal firm-size choice across idea qualities in inventor-led innovation. A positive firm size corresponds to joining an incumbent, whereas zero indicates founding a new business. The figure indicates positive assortative matching among inventors who join incumbents: higher-quality ideas are matched with larger firms. Once idea quality, z_0, exceeds a cutoff, the inventor instead chooses to launch a startup. This pattern is consistent with the fact that patents produced by startups receive more citations on average than those produced by incumbents in our data.

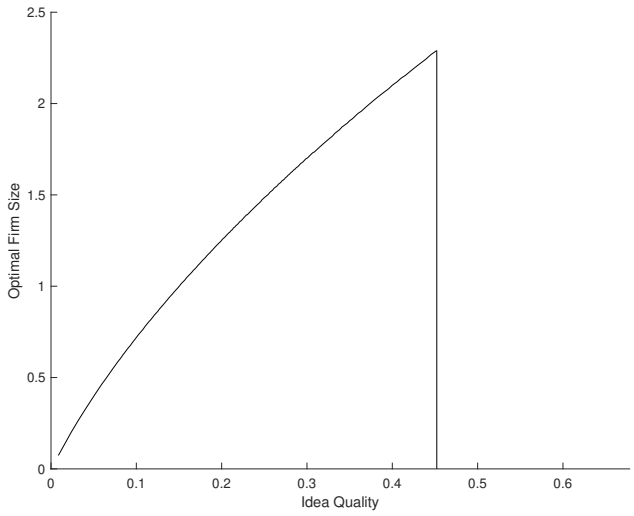


Figure 5: Firm-size Choices of Inventors

Notes: The figure shows the model-implied mapping from an inventor's idea quality to the optimal firm size. A positive firm size corresponds to joining an incumbent of that size, whereas zero indicates founding a new business.

Variation in technology waves is modeled as changes in the scale parameter of the idea-quality distribution, z_m. A higher z_m implies better and more frontier-shifting ideas on average. As z_m rises, more inventors choose to found startups and positive

assortative matching between idea quality and firm size weakens. Both forces reduce market concentration. Moreover, with more ideas developed in startups, growth shifts toward expansion in the mass of firms rather than increases in average firm quality.

5 Quantitative Analysis

Using the calibrated model, we simulate the economy starting in 1986, the first peak of the technology wave in our sample. Appendix D.4 reports analogous simulations starting in 1984, confirming that our results are robust to the choice of initial year. In each year, we adjust the scale parameter of the idea quality distribution, z_m, so that the model-implied Tech Wave Index, $\frac{\int z_0 \, d\Psi(z_0)}{B+\int z_0 \, d\Psi(z_0)}$, matches its empirical counterpart, holding all other parameters fixed. The simulation produces annual series for model-implied growth rates, market concentration, and the allocation of innovation. We compare these series to the data and compute their correlations. Appendix D.5 reports simulation results in which the model-implied Tech Wave Index is matched to the "RETech" index developed by Bowen III, Frésard and Hoberg (2023). The close alignment with the baseline results demonstrates robustness to alternative measures of technology waves.

The initial firm-size distribution is pinned down by the empirical distribution from the Business Dynamics Statistics (BDS). Appendix D.1 describes the procedure in detail. Because the equilibrium features no transition dynamics, the evolution of firm mass $N_F(t)$ and the firm-size distribution $f_t(q)$ can be computed directly. The resulting paths of $N_F(t)$ and $f_t(q)$ jointly determine the time series of growth rates and market concentration.

5.1 Growth Rate

The aggregate growth rate, g, can be decomposed into two components: growth in average firm quality and growth in firm mass (Equation (38)). Using the simulated economy, we compute 5-year averages of g and of each component, and report them in Figure 6, alongside the HP-filtered annual growth rate averaged over 5-year windows in the data.[30] Two patterns stand out. First, aggregate growth declines after 2000, with most of the slowdown coming from the decline in firm mass. Second, average firm quality behaves differently: growth in average firm quality is relatively strong in 2000–2010, consistent with faster improvements among surviving firms, but it is more than offset by a pronounced contraction in firm mass. Aggregate growth, g, rebounds in the last five years, reflecting positive growth in both average quality and firm mass, although it

[30]The model-generated annual growth rate and its empirical counterpart are shown in Figure 19 in Appendix D.2.

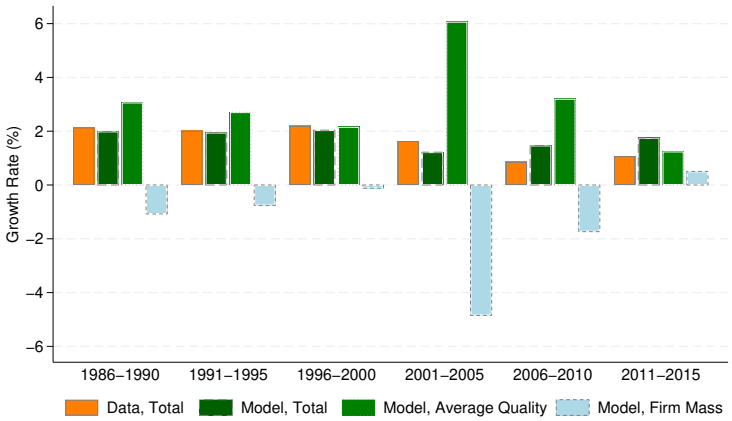


Figure 6: Decomposition of Growth Rates by Period

Notes: This figure reports 5-year averages of the aggregate economic growth rate and its two components. In each 5-year window, the left bar shows total growth. The middle bar shows the contribution from improvements in average firm quality, and the right bar shows the contribution from changes in firm mass.

does not return to the pre-2000 level. These qualitative patterns align with the evidence in Autor et al. (2020) and Akcigit and Ates (2023): since the early 2000s, productive incumbents have expanded while net firm entry and overall dynamism have declined.

5.2 Market Concentration

Figure 7 presents the simulated evolution of market concentration, measured by the Herfindahl–Hirschman Index (HHI), together with its empirical counterparts. Appendix D.3 details the construction of the model-simulated HHI. The solid red line shows the model-simulated HHI, the dashed purple line reports the employment-based HHI from Census data, and the dash-dotted pink line reports the sales-based HHI from Compustat. All series are normalized by their mean over 1986–2016. The model-generated HHI closely tracks both concentration measures in the data.

Although the calibration does not explicitly target any measure of market concentration, the model successfully replicates both the upward trend and cyclical fluctuations observed in the data. To disentangle these components, we fit separate linear trends to the simulated and empirical HHI series and subtract them to obtain the detrended time variations, as shown in Figure 8. Summary statistics are reported in the first three rows of Panel A in Table 9.

The fitted linear trend has a slope of 0.011, close to the trends in both the employment- and sales-based empirical HHI measures. The standard deviation of the model's

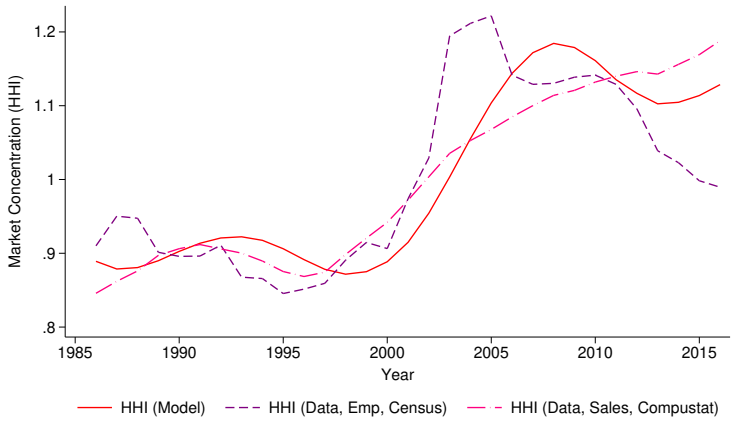


Figure 7: Model-generated HHI and Empirical Counterparts

Notes: This figure compares model-generated and empirical market concentration over time. The solid red line shows the simulated annual HHI, normalized by its mean over 1986–2016. The dashed purple line reports the employment-based HHI from the LBD, and the dash-dotted pink line reports the sales-based HHI from Compustat; both are normalized by their respective means.

detrended series is 0.060, which lies between the corresponding values for the two empirical measures. The model's first-order autocorrelation is 0.916, compared with 0.847 in the employment-based data and 0.930 in the sales-based data. The last column reports the correlation between the detrended model-generated HHI and the detrended empirical HHI series: 0.658 for the employment-based measure and 0.793 for the sales-based measure (in parentheses). Taken together, these results indicate that technology waves are an important driver of both the long-run trend and cyclical fluctuations in market concentration.

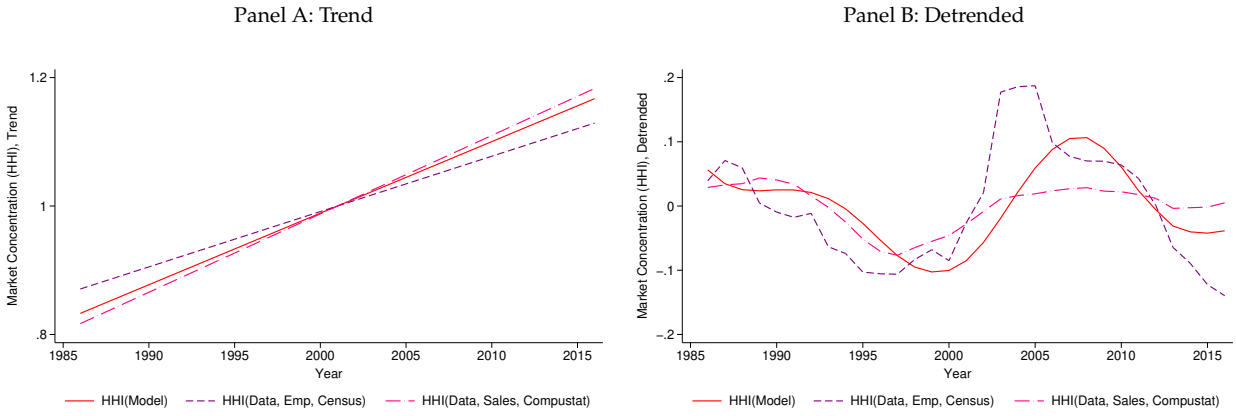


Figure 8: The Trend and Detrended Time Variations of the HHI

Notes: The left panel plots the linear trend in market concentration (HHI) over time in the model (solid red line) and in the data (dashed purple line for employment and dash-dotted pink link for sales). The right panel shows the corresponding detrended HHI series, highlighting fluctuations around the trend.

Table 9: Comparison between Model and Data

	Original		Detrended		
	Mean	Slope	S.D.	Autocorr	Corr. Data
			Panel A. HHI		
Data (Emp, Census)	1	0.009	0.092	0.847	1 (0.621)
Data (Sales, Compustat)	1	0.012	0.035	0.930	0.621 (1)
Model	1	0.011	0.060	0.916	0.658 (0.793)
Model (Counterfactual)	1.10	0.014	0.053	0.918	0.660 (0.767)
			Panel B. New-to-Incumbent Ratio		
Data	0.052	-5.05E-4	0.007	0.830	1
Model	0.055	-8.27E-4	0.012	0.910	0.763

Notes: This table summarizes the trend and detrended time-series variation in the HHI and the new-to-incumbent ratio in both the data and the model. The columns report the series mean, the slope of the fitted linear trend, the standard deviation of the detrended series, the first-order autocorrelation of the detrended series, and the correlation of the detrended series with its data counterpart. For the HHI, two correlations are reported: with the employment-based measure and, in parentheses, with the sales-based measure.

5.3 Decomposition of the Extensive and Intensive Margins

Market concentration is shaped by both the extensive margin (the mass of firms) and the intensive margin (matching between idea quality and firm size). We use the allocation of ideas and a counterfactual exercise to examine the contribution of each margin.

5.3.1 Extensive Margin: Allocation of Ideas

Idea allocation shapes the evolution of firm mass. Empirically, inventors are more likely to found startups when breakthrough technologies emerge and to join incumbent firms as technologies mature. Panel A of Figure 9 shows that the empirical new-to-incumbent ratio (dashed) and the model-implied series (solid) display closely aligned dynamics. After removing separate linear trends, the two series continue to co-move strongly, with a detrended correlation of 0.763. These results indicate that technology waves account for the bulk of the evolution in firm mass.

Panel B of Table 9 reports additional summary statistics. The average New-to-Incumbent ratio in the model is 0.055, close to 0.052 in the data. The slope of the linear trend in the model is $-8.27E-3$, slightly more negative than the data counterpart. The detrended series also shows slightly larger variability, with a standard deviation of 0.012 compared to 0.007 in the data, and a slightly higher first-order autocorrelation (0.910 vs. 0.830). The model's larger volatility may be due to the absence of adjustment costs. In



Figure 9: Idea Allocation and the Effect of PAM

Notes: Panel A reports idea allocation over time: the solid line plots the model-simulated new-to-incumbent ratio, and the dashed line plots its empirical counterpart (as in Figure 3). Panel A illustrates the extensive-margin component of model-implied market concentration over time. The solid line shows the simulated HHI from the full model, while the dashed line shows the counterfactual HHI driven solely by the extensive margin. Both series are normalized by their respective mean values.

our model, inventors are short-lived and choose between joining startups or incumbent firms without regard to their previous affiliations. As a result, the New-to-Incumbent ratio reacts immediately to aggregate shocks.

5.3.2 Intensive Margin: Positive Assortative Matching

To isolate the intensive-margin effect, we consider a counterfactual in which inventors who choose startup-led innovation are matched to firm size through the same assortative mechanism as inventors joining incumbents. To shut down the extensive-margin channel, inventors first make the same startup-versus-incumbent choice as in the baseline, with startup size still drawn from the entry distribution, $f_{I,t}(\tilde{q}_I)$. Conditional on founding a startup, the inventor also solves the same startup contracting problem as in the baseline. However, her idea is then assigned to a firm of the optimal size implied by its quality, following the same assortative matching rule as for incumbents. This exercise leaves firm mass evolution the same as the baseline and alters only the matching between idea quality and firm size.

We re-simulate the economy under this counterfactual. Panel B of Figure 9 plots the resulting HHI (dashed) against the baseline (solid), with both series normalized by the mean baseline HHI, and Panel A of Table 9 reports the corresponding moments. In the counterfactual, average HHI is 10% higher, and its upward trend is steeper, with a slope of 0.14 rather than 0.11. The reason is that the attenuation of positive assortative matching in startup formation is removed. This shows that variation in the intensive margin over

technology waves contributes meaningfully to the market concentration evolution.

5.4 Startup tax

The literature has documented increasing barriers to entry and growing patent thickets, especially since the 2000s (Akcigit and Ates, 2023; Hall, Graevenitz and Helmers, 2021). Besides, entry by new businesses has become increasingly dependent on computing resources and cloud services provided by incumbent firms.[31] To capture these changes in a reduced-form way, we introduce a proportional startup tax $\xi \geq 0$ that applies only when an inventor chooses to start a new business. This tax can be interpreted as either a stricter effective patent policy (e.g., higher IP-related transaction costs in thicketed areas) or payment for essential goods or services provided by incumbent firms (e.g., cloud compute and complementary services).

Under the tax, the startup effort cost is increased by a multiplicative wedge:

$$R(e_I; \xi) \equiv (1 + \xi)\, R(e_I),$$

while all non-startup inventor problems are unchanged. We define the associated tax revenue as the wedge ξ applied to the (pre-tax) startup effort cost flow:

$$\mathcal{T}(t) \equiv h\xi\, V(t \mid Q(t)) \int_{z_0 \in \{z_0 | \tilde{q}^* = 0\}} R\big(e_I(z_0, 0, t)\big)\, d\Psi(z_0),$$

where $\tilde{q}^* = 0$ indexes inventors who optimally choose the startup option.

Assume these tax revenues are rebated to incumbent intermediate firms as a revenue subsidy. Let $\zeta(t) \geq 0$ denote the subsidy rate applied to firm j's revenue, so firm j solves

$$\max_{p_j(t),\, y_j(t),\, \ell_j(t)} (1 + \zeta(t))\, p_j(t)\, y_j(t) - w(t)\, \ell_j(t).$$

Period-by-period balance is imposed so that total tax revenues equal total rebates to the intermediate sector:

$$\mathcal{T}(t) \;=\; \zeta(t) \int_0^{N_F(t)} p_j(t)\, y_j(t)\, dj \;=\; \zeta(t) \int_0^{N_F(t)} q_j(t)\, dj \;=\; \zeta(t)\, Q(t).$$

We conduct a simulation in which the economy starts from the 1986 equilibrium with

[31] Recent U.S. policy evidence highlights that leading cloud providers increasingly partner with frontier AI developers through arrangements that bundle large investments with substantial cloud spending commitments ("circular spending"), preferential access to scarce compute, and other contractual and technical features (including data-transfer/egress charges) that can raise switching costs and deepen reliance on incumbent cloud infrastructure (Federal Trade Commission, 2025).



Figure 10: Simulated HHI and New-to-Incumbent Ratio Under a Startup Tax

Notes: Panel A illustrates the simulated HHI. The solid line shows the baseline path, while the dashed and dash-dotted lines report the time series under $\xi_T = 5\%$ and $\xi_T = 10\%$, respectively. Both tax scenarios are normalized by the baseline mean HHI. Panel B plots the idea new-to-incumbent ratio over time: baseline (solid), 5% tax (dashed), and 10% tax (dash-dotted).

zero startup wedge, $\xi_0 = 0$, and converges to a terminal value $\xi_T \in \{0.05, 0.10\}$ by 2016 over a 30-year transition ($T = 30$). Following Akcigit and Ates (2023), we parameterize the startup wedge along the transition using the following functional form,

$$\xi_t = \xi_0 + f(t; \nu_\xi)(\xi_T - \xi_0), \qquad f(t; \nu_\xi) \equiv \frac{\exp\left(\frac{t}{T}\nu_\xi\right) - 1}{\exp(\nu_\xi) - 1} \in [0, 1].$$

We set the curvature (speed) parameter to $\nu_\xi = 1$, so that ξ_t rises smoothly from its initial to its terminal level. Given the exogenous sequence $\{\xi_t\}_{t=0}^{T}$, we re-simulate the economy and report the implied time series of the HHI and the new-to-incumbent ratio under $\xi_T = 5\%$ and 10%. Figure 10 compares these paths to the baseline. The HHI series under the tax scenarios are normalized by the baseline mean HHI.

Under startup taxes, the markets become more concentrated compared to the baseline, especially post-2000s. In Panel A, the three series track each other closely before 2000, but the tax scenarios start to diverge in the early 2000s and display a noticeably steeper increase in concentration through the 2000s. The long period lack of technological breakthrough between the mid-1990s and early 2010s pushes the HHI to a higher level, and the later breakthroughs has smaller effect in lowering concentration. The New-to-Incumbent ratio shown in Panel B is lower, especially after the 2000s, with more ideas flowing to incumbent firms. Panel B also shows that while the ratio partially rebounds around the early 2010s, the recovery is weaker under startup taxes and the gap relative to the baseline widens toward 2016, consistent with a persistent shift of innovative output away from entrants and toward incumbents.

6 Conclusion

This paper studies how technology waves shape growth and market concentration through the reallocation of ideas. Combining empirical evidence with a structural analysis, we show that market concentration is inversely related to technology waves, pointing to a low-frequency business-cycle component in the economy. Using the Longitudinal Business Database (LBD) from the U.S. Census Bureau and patent data from the USPTO, we document that the share of patents developed in new businesses comoves closely with technology waves: at wave peaks, a larger share of patents originates in new businesses, whereas at troughs, patenting shifts toward incumbents.

This paper proposes a theoretical framework that elucidates the decision-making process of inventors regarding their choice of innovation pathways, thus providing an explanation for the observed empirical patterns. Our model effectively captures the contraction in firm mass and the increase in average firm quality after the 2000s, and both the trend and fluctuations in market concentration. It implies that the deceleration in the emergence of groundbreaking technologies could be a significant contributing factor to the decline in economic growth and the rise in market concentration after the 2000s.

The rise of generative AI may reduce market concentration by encouraging more startup formation. A surge in AI business applications has already been observed in the data (Dinlersoz, Dogan and Zolas (2024)). However, the magnitude of this effect depends on how much new firms rely on incumbents' platforms and products. If these dependencies are strong, incumbents may capture a large share of the rents from AI-driven entry, dampening the impact on market concentration.

References

Acemoglu, Daron, Ufuk Akcigit, and William R Kerr. 2016. "Innovation network." *Proceedings of the National Academy of Sciences*, 113(41): 11483–11488.

Aghion, Philippe, and Jean Tirole. 1994. "The management of innovation." *The Quarterly Journal of Economics*, 109(4): 1185–1209.

Aghion, Philippe, and Peter Howitt. 1990. "A model of growth through creative destruction."

Akcigit, Ufuk, and Nathan Goldschlag. 2023. "Where have all the" creative talents" gone? Employment dynamics of us inventors." National Bureau of Economic Research.

Akcigit, Ufuk, and Sina T Ates. 2023. "What happened to US business dynamism?" *Journal of Political Economy*, 131(8): 2059–2124.

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

Akcigit, Ufuk, Murat Alp Celik, and Jeremy Greenwood. 2016. "Buy, Keep, or Sell: Economic Growth and the Market for Ideas." *Econometrica*, 84(3): 943–984.

Atkeson, Andrew, and Patrick J Kehoe. 2007. "Modeling the transition to a new economy: lessons from two technological revolutions." *American Economic Review*, 97(1): 64–88.

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

Ayerst, Stephen, Faisal Ibrahim, Gaelan MacKenzie, and Swapnika Rachapalli. 2023. "Trade and diffusion of embodied technology: an empirical analysis." *Journal of Monetary Economics*, 137: 128–145.

Bena, Jan, and Kai Li. 2014. "Corporate Innovations and Mergers and Acquisitions." *The Journal of Finance*, 69(5): 1923–1960.

Bighelli, Tommaso, Filippo Di Mauro, Marc J Melitz, and Matthias Mertens. 2023. "European firm concentration and aggregate productivity." *Journal of the European Economic Association*, 21(2): 455–483.

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb. 2020. "Are ideas getting harder to find?" *American Economic Review*, 110(4): 1104–1144.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-experimental shift-share research designs." *The Review of economic studies*, 89(1): 181–213.

Bowen III, Donald E, Laurent Frésard, and Gerard Hoberg. 2023. "Rapidly evolving technologies and startup exits." *Management Science*, 69(2): 940–967.

Campbell, Jeffrey R, and Hugo A Hopenhayn. 2005. "Market size matters." *The Journal of Industrial Economics*, 53(1): 1–25.

Cassiman, Bruno, and Reinhilde Veugelers. 2006. "In Search of Complementarity in Innovation Strategy: Internal R&d and External Knowledge Acquisition." *Management Science*, 52(1): 68–82.

Cavenaile, Laurent, Murat Alp Celik, and Xu Tian. 2019. "Are markups too high? competition, strategic innovation, and industry dynamics." *Competition, Strategic Innovation, and Industry Dynamics (August 1, 2019)*.

Chatterjee, Satyajit, and Esteban Rossi-Hansberg. 2012. "Spinoffs and the Market for Ideas." *International Economic Review*, 53(1): 53–93.

Cunningham, Colleen, Florian Ederer, and Song Ma. 2021. "Killer acquisitions." *Journal of Political Economy*, 129(3): 649–702.

Dinlersoz, Emin, Can Dogan, and Nikolas Zolas. 2024. "Starting Up AI."

Federal Trade Commission. 2025. "Partnerships Between Cloud Service Providers and AI Developers (6(b) Study Report)." Federal Trade Commission.

Fons-Rosen, Christian, Pau Roldan-Blanco, and Tom Schmitz. 2021. "The Aggregate Effects of Acquisitions on Innovation and Economic Growth."

Ganapati, Sharat. 2021. "Growing oligopolies, prices, output, and productivity." *American Economic Journal: Microeconomics*, 13(3): 309–327.

Giarratana, Marco S, Myriam Mariani, and Ingo Weller. 2018. "Rewards for patents and inventor behaviors in industrial research and development." *Academy of Management Journal*, 61(1): 264–292.

Greenwood, Jeremy, and Mehmet Yorukoglu. 1997. "1974." Vol. 46, 49–95, Elsevier.

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

Grullon, Gustavo, Yelena Larkin, and Roni Michaely. 2019. "Are US industries becoming more concentrated?" *Review of Finance*, 23(4): 697–743.

Hall, Bronwyn, and John Van Reenen. 2000. "How effective are fiscal incentives for R&D? A review of the evidence." *Research policy*, 29(4-5): 449–469.

Hall, Bronwyn H, Georg von Graevenitz, and Christian Helmers. 2021. "Technology entry in the presence of patent thickets." *Oxford Economic Papers*, 73(2): 903–926.

Holmstrom, Bengt. 1989. "Agency costs and innovation." *Journal of Economic Behavior & Organization*, 12(3): 305–327.

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

Jovanovic, Boyan, and Peter L Rousseau. 2014. "Extensive and intensive investment over the business cycle." *Journal of Political economy*, 122(4): 863–908.

Kalyani, Aakash. 2024. "The Creativity Decline: Evidence from US Patents." *Available at SSRN 4318158*.

Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, and Ahmed Tahoun. 2025. "The diffusion of new technologies." *The Quarterly Journal of Economics*, 140(2): 1299–1365.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring technological innovation over the long run." *American Economic Review: Insights*, 3(3): 303–320.

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

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

Kwon, Spencer Yongwook, Yueran Ma, and Kaspar Zimmermann. 2023. "100 years of rising corporate concentration." *University of Chicago, Becker Friedman Institute for Economics Working Paper*, , (2023-20).

Kydland, Finn E, and Edward C Prescott. 1982. "Time to build and aggregate fluctuations." *Econometrica: Journal of the Econometric Society*, 1345–1370.

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

Lerner, Josh, and Julie Wulf. 2007. "Innovation and incentives: Evidence from corporate R&D." *the Review of Economics and Statistics*, 89(4): 634–644.

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

Liu, Huben, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller. 2025. "Technology and Labor Markets: Past, Present, and Future; Evidence from Two Centuries of Innovation." National Bureau of Economic Research.

Ma, Yueran, Mengdi Zhang, and Kaspar Zimmermann. 2024. "Business Concentration around the World: 1900—2020."

Ma, Yueyuan. 2026. "Specialization in a knowledge economy." *Journal of Political Economy Macroeconomics*, 4(1): 000–000.

Merges, Robert P. 1999. "The law and economics of employee inventions." *Harv. JL & Tech.*, 13: 1.

Olmstead-Rumsey, Jane. 2019. "Market concentration and the productivity slowdown."

Onishi, Koichiro. 2013. "The effects of compensation plans for employee inventions on R&D productivity: New evidence from Japanese panel data." *Research Policy*, 42(2): 367–378.

Perla, Jesse, Christopher Tonetti, and Michael E. Waugh. 2021. "Equilibrium Technology Diffusion, Trade, and Growth." *American Economic Review*, 111(1): 73–128.

Romer, Paul M. 1986. "Increasing Returns and Long-Run Growth." *Journal of Political Economy*, 94(5): 1002–1037.

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy*, 98(5, Part 2): S71–S102.

Yang, Shaoshuang. 2026. "The Macroeconomic Implications of R&D Ownership." *Working Paper*.

Online Appendix

A Data Description

The data used in this paper includes the Longitudinal Business Database (LBD), the USPTO patent data, the Compustat Fundamentals Annual, and PATSTAT Global (Spring 2024). This section provides details about the information of the datasets and the construction of key variables.

A.1 The USPTO Patent Data

The USPTO patent data contains information of all patents issued between 1976 and 2022. It can be downloaded from the PatentsView website. For each patent, the data documents the patent type (utility, design, plant, etc.), the IPC code indicating its technological class, the grant year, and the patents it cites and it is cited. We keep all the utility patents to focus our attention to the introduction of new products and processes.

Forward Citations Forward citations are citations a focal patent receives from others. It indicates how many patents follow the focal one. This paper calculates the number of forward citations each patent gets within five years after issuance.

Backward Citations Backward citations are citations that other patents receive from the focal patent. It indicates to what extent the focal patent follows the existing technology. This paper calculates the number of backward citations by counting the number of patents cited by the focal patent that were granted within the previous five years.

The Tech Wave Index According to the definition in the paper, we calculate this index by dividing the number of forward citations received by all the utility patents granted in a year by the summation of the forward and backward citations of those patents. The Tech Wave Index by IPC is derived in a similar way for each IPC class and each year.

A.2 The Compustat Fundamentals Annual

The Compustat Fundamentals Annual contains information of all the publicly listed firms in the US. It records the firms' net sales, the number of employees, the primary industry (4-digit SIC code), and the headquarter locations of each firm. We keep all the firms that are headquartered in the US.

Primary Industry The primary industry of each firm in Compustat is based on the 4-digit SIC code assigned to each firm in the Fundamentals Annual. The code can be aggregated to different levels. Manufacturing is corresponding to SIC codes 2000-3999; utility and

transportation is corresponding to SIC codes 4000-4999; wholesale trade is corresponding to SIC codes 5000-5199; retail trade is corresponding to SIC codes 5200-5999; finance is corresponding to SIC codes 6000-6999; service is corresponding to SIC codes 7000-8999.
The Herfindahl-Hirschman Index (HHI) Following the methods in Grullon, Larkin and Michaely (2019), we first calculate the HHI of each 3-digit SIC code by the squared ratios of firm net sales to the total net sales in that 3-digit industry. To get the aggregate HHI, we sum up the HHIs of all the 3-digit SIC codes and weight them by their total net sales.

A.3 The Longitudinal Business Database (LBD)

The LBD is collected by the US Census Bureau and is an establishment-level data that covers the universe of US businesses with paid employees from 1976 to 2020. The dataset assigns a firm ID to all establishments belonging to the same firm. Using the Business Dynamics Statistics of Patenting Firms (BDS-PF) patent assignee-FIRMID crosswalk from the Census, this paper links the USPTO patent data with firms in the LBD, therefore, derives all utility patents in the US that were granted to employer businesses between 1976 and 2020.
New-to-Incumbent Ratio After merging the patent data with the LBD, this paper can identify the firm each patent was granted to. If the firm is less than or equal to five years old in the patent's grant year, we indicate that the idea behind the patent was absorbed by a new firm 5 years ago. Otherwise, we indicate that the idea was absorbed by an incumbent firm 5 year ago. Then we divide the number of ideas combined with new firms by the number of ideas combined with incumbent firms to get the New-to-Incumbent Ratio.
Firm Size The LBD documents the number of employees each firm hires in each year. We deriving the mapping between patent forward citations and incumbent firm size, we use the number of employees as a proxy for size.

A.4 PATSTAT (Patent Statistical Database)

PATSTAT Global (Spring 2024) is a comprehensive global dataset maintained by the European Patent Office (EPO). PATSTAT provides detailed bibliographic data on patents from various patent offices worldwide, with a particular focus on those filed through the EPO. We restrict our sample to patents with inventors based in European countries.
Forward citations, backward citations, and the Tech Wave Index for European patents are constructed in the same way as for U.S. patents using USPTO data.

B More Empirical Evidence

B.1 Empirical Patterns without Smoothing

Figure 11 and Figure 12 present the patterns of the technology waves, HHI, and the New-to-Incumbent ratio without smoothing techniques. The negative correlation between the HHI and the Tech Wave Index, as well as the positive correlation between the New-to-Incumbent ratio and the Tech Wave Index, remain prominent.

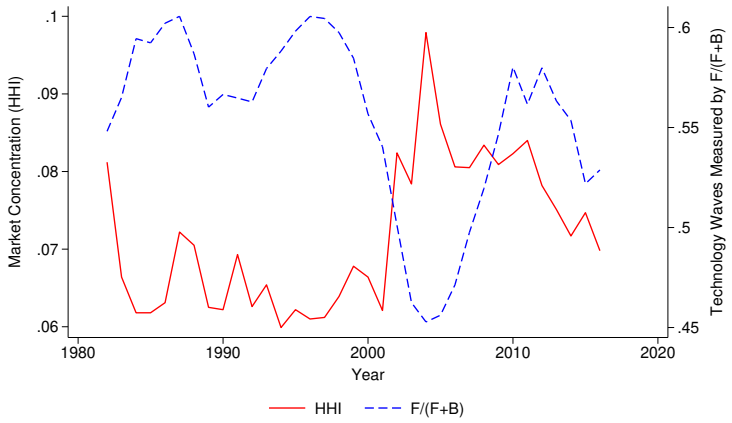


Figure 11: Technology Waves and Market Concentration without Smoothing

Notes: This figure shows the technology waves and market concentration of employment over time. The blue dashed curve, based on the methodology defined in this paper, calculates the relative ratio of forward citations to the sum of forward and backward citations. The red solid curve displays the HHI in each year, which is the weighted average of the industry-level HHI in each year. The weight is the total employment of firms in each industry. The two curves have different y-axises, shown respective on the left and right.
Sources: Compustat Fundamental Annuals and USPTO patent and citation data.

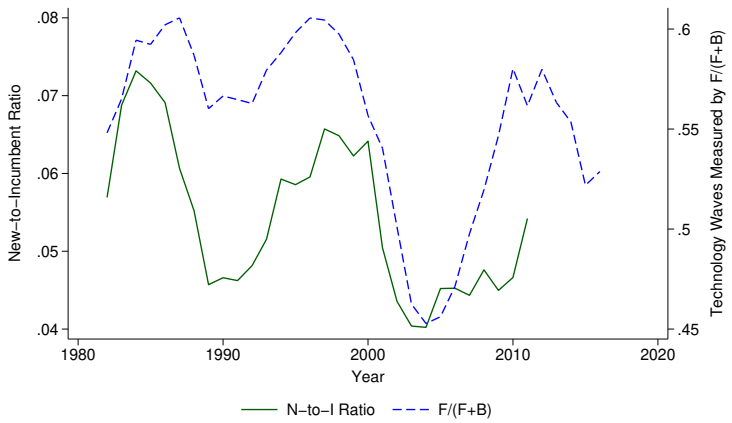


Figure 12: Technology Waves and Idea Allocation without Smoothing

Notes: This figure shows the technology waves and the idea allocation between new and incumbent firms over time. The blue dashed curve, based on the methodology defined in this paper, calculates the relative ratio of forward citations to the sum of forward and backward citations. The green solid curve displays the "New-to-Incumbent Ratio" defined in the paper, capture where new ideas contribute their value. The two curves have different y-axises, which are shown respective on the left and right.
Sources: Longitudinal Business Database (LBD) and USPTO patent and citation data.

B.2 Technology waves by Technological Field

The Tech Wave Index across the nine technological fields is shown in Figure 13. The index is based on the same algorithm as in Equation (1) except that the forward and backward citations are aggregated across each of the 1-digit IPC code. The top three fields with the highest Tech Wave Index are Human Necessities, Physics, and Electricity at the first peak; Electricity, Physics, and Human Necessities at the second peak; Human Necessities, Chemistry and Metallurgy, and Mechanical Engineering etc. at the third peak.

B.3 Main Drivers of the Wave Peaks

To capture the main technology classes that drive the peaks of the technology waves, we decompose the Tech Wave Index into the contribution of each three-digit IPC code using the following method,

$$\text{TWI}_t = \sum_{j \in J} \frac{\sum_{i \in I_{jt}} \sum_{s=0}^{5} F_{ij,t+s}}{\sum_{i \in I_{jt}} \sum_{s=0}^{5} F_{ij,t+s} + \sum_{i \in I_{jt}} \sum_{s=0}^{5} B_{ij,t-s}} \frac{\sum_{i \in I_{jt}} \sum_{s=0}^{5} F_{ij,t+s} + \sum_{i \in I_{jt}} \sum_{s=0}^{5} B_{ij,t-s}}{\sum_{i \in I_t} \sum_{s=0}^{5} F_{i,t+s} + \sum_{i \in I_t} \sum_{s=0}^{5} B_{i,t-s}}, \tag{41}$$

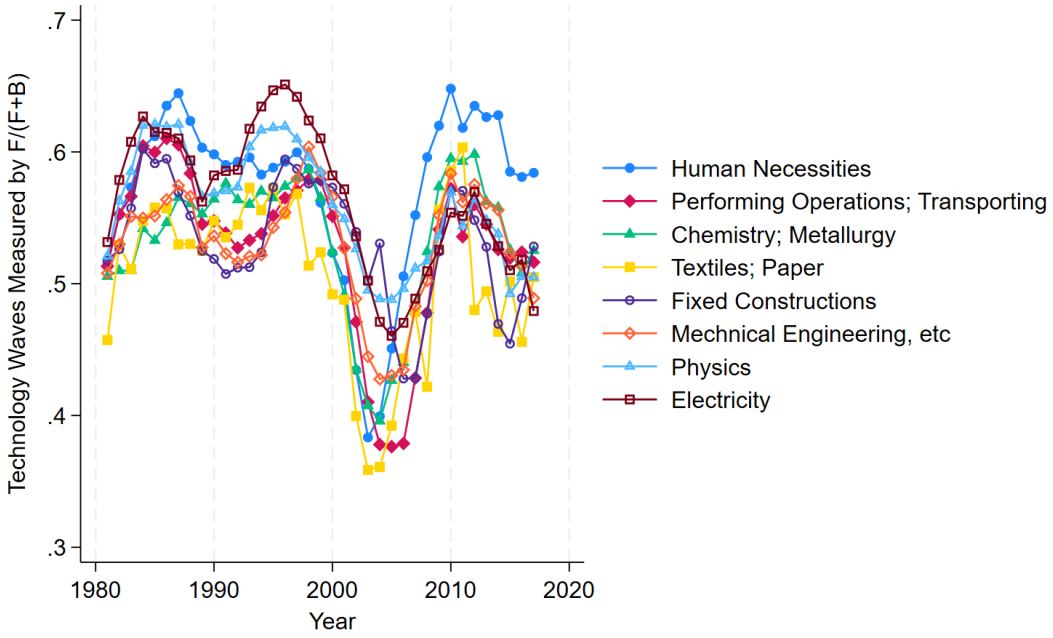


Figure 13: Technology Waves by Technological Fields

Notes: This figure shows the technology waves by the nine technological fields between 1981 and 2017. The nine fields are defined by the 1-digit IPC code. The technology waves are measured by the Tech Wave Index as defined by Equation 1 in the paper.
Sources: USPTO patent and citation data.

where J is the set of 3-digit IPC code and I_{jt} is the set of patents belonging to the IPC code j granted in year t. Intuitively, the contribution of each technology class in a given year is determined by the IPC-specific Tech Wave Index multiplied by the share of forward and backward citations of that class. Table 10 lists the top three contributors at the three peaks of the technology waves. Medical or Veterinary Science and Hygiene contribute most to the first peak, while Computing; Calculating or Counting is the leading contributor to the second and third peak.

Table 10: Major Contributors to the Peaks of Technology Waves

	First Peak (1985-1987)	Second Peak (1995-1997)	Third Peak (2010-2012)
1	Medical or Vet. Sci.; Hygiene	Computing; Calculating or Counting	Computing; Calculating or Counting
2	Electric Elements	Medical or Vet. Sci.; Hygiene	Medical or Vet. Sci.; Hygiene
3	Measuring; Testing	Electric Communication Technique	Electric Communication Technique

Notes: This table shows the major technological classes of the top three fields with the highest Tech Wave Index at the peaks of technology waves in the period between 1981 and 2017.

B.4 Applicants of Breakthrough Patents

To examine the second question, we draw on the "historically significant patents" compiled by Kelly et al. (2021) from online sources. Our sample period includes 54 such breakthrough patents. To assess whether these patents meaningfully contribute to aggregate technology waves, we compute a patent-level Tech Wave Index using the same methodology applied to the IPC-level index.[32] We then rank each of the 54 patents within the unconditional distribution of Tech Wave Index scores over the sample period. The mean and median percentile ranks fall in the top 24% and top 7% of all patents, respectively, indicating a strong concordance between Kelly et al. (2021)'s list and our novelty measure. We categorize the applicants of these breakthrough patents by institution type and find that 41% are incumbent firms, 26% are startups, 20% are universities, and the remainder come from other institutions, as shown in Figure 14. The wide range of sources demonstrates that frontier-shifting technologies emerge across diverse organizational settings instead of a dominant sector.

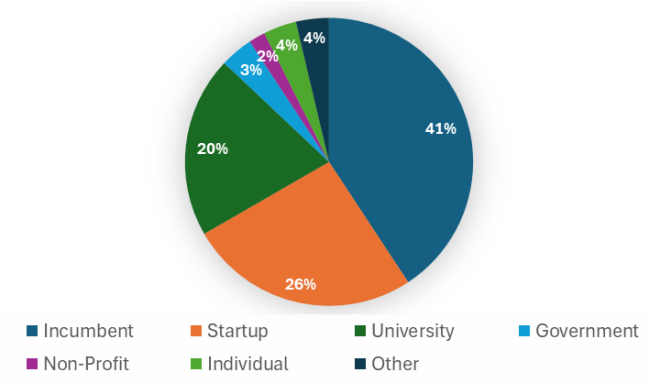


Figure 14: Composition of Applicants of Breakthrough Patents

Notes: The pie chart illustrates the share of applicants from each institutional group for the 54 breakthrough patents. Incumbents are defined as private firms that applied for a patent at least three years after their founding. Startups are private firms that applied within their first three years or before. Others include cases where the applicants are multiple institutions of different types.
Sources: The "historically significant patents" complied by Kelly et al. (2021).

B.5 Tech Wave Index in Europe

To calculate the Tech Wave Index for European countries with intensive patenting activities, we use data from PATSTAT Global (Spring 2024). The six countries with the

[32] Namely,

$$\text{TWI}_t = \sum_{i \in I_t} \frac{\sum_{s=0}^{5} F_{i,t+s}}{\sum_{i \in I_t} \sum_{s=0}^{5} F_{i,t+s}} \times \frac{\sum_{s=0}^{5} F_{i,t+s}}{\sum_{i \in I_t} \sum_{s=0}^{5} F_{i,t+s} + \sum_{i \in I_t} \sum_{s=0}^{5} B_{i,t-s}}.$$

highest number of patent issuances between 1982 and 2016 are Germany, France, the United Kingdom, Italy, Switzerland, and the Netherlands. Using the definition of the Tech Wave Index as outlined in Equation (1), we calculate the technology waves in these six countries and present them in Figure 15. Across all six countries, we observe an overall declining trend in technological novelty, with common peaks in the mid-1980s and early 2010s. Italy also experienced a distinct peak in the mid-1990s. In general, the technological trends in these European countries with the highest patenting activity mirror those observed in the U.S.[33]

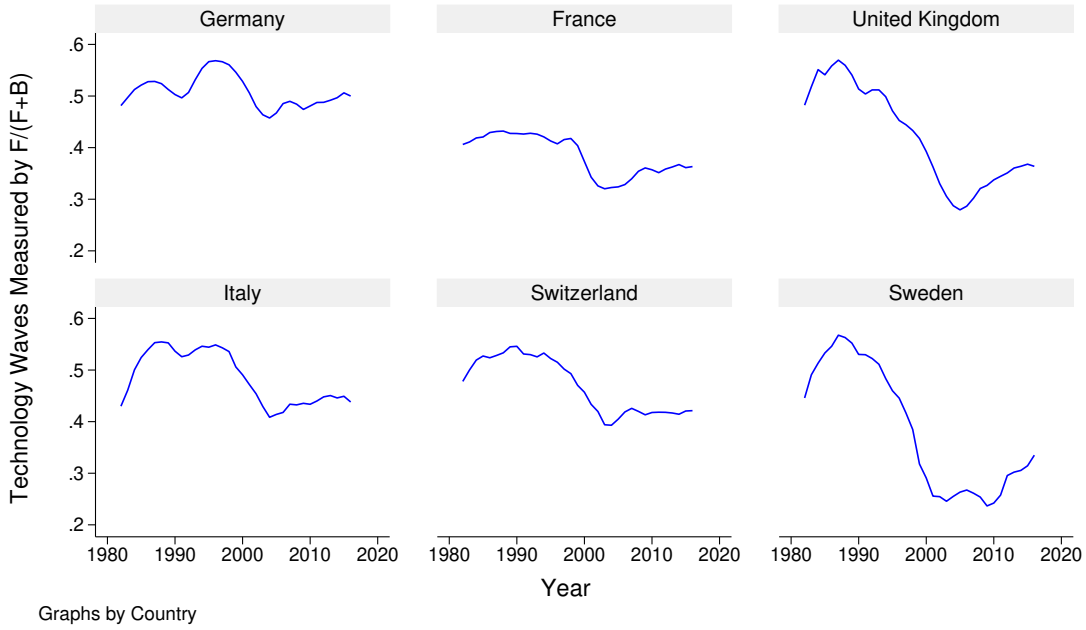


Figure 15: technology waves in European Countries

Notes: This figure shows the technology waves in six European countries with the highest number of patent issuances between 1982 and 2016. The technology waves are measured by the Tech Wave Index as defined by Equation 1 in the paper.
Sources: PATSTAT (Patent Statistical Database).

B.6 Relationship with the Technology Waves

Table 11 exhibits the time trend of the technological novelty waves, the market concentration measured by the HHI, and the New-to-Incumbent Ratio of idea allocations (Panel A). It also displays the cross correlation of the two latter time series with the technology waves at different year gaps (Panel B). The time trend is derived by fitting

[33]Similar results can be obtained by calculating the Tech Wave Index using Google Patents data, as cleaned by Ayerst et al. (2023). These results are available upon request.

a linear trend to the focal time series and taking its slope. The cross correlations are obtained by calculating the correlation coefficients of the detrended time series when the year gaps of the two series are respectively -3, -2, -1, 0, 1, 2, 3. The detrending process subtracts the linear trend from the original time series. The cross correlations capture not only the co-movement of the different time series, but also the relative timing of their movements. The first row of each panel shows the statistics for the whole sample; the subsequent rows are statistics by major industries according to the Standard Industrial Classification (SIC) code or technological fields according to the International Patent Classification (IPC).

Table 11: Time Trend and Cross Correlation

	Time Trend		Detrended Cross Correlation						
			Panel A. HHI						
	Tech Wave	HHI	$k=-3$	$k=-2$	$k=-1$	$k=0$	$k=1$	$k=2$	$k=3$
All	-0.002	0.007	-0.625	-0.720	-0.775	-0.736	-0.588	-0.366	-0.097
			Panel B. New-to-Incumbent Ratio						
	Tech Wave	N-to-I Ratio	$k=-3$	$k=-2$	$k=-1$	$k=0$	$k=1$	$k=2$	$k=3$
All	-0.002	-0.001	0.107	0.314	0.504	0.612	0.536	0.317	-0.001
Human Necessities	-0.001	-0.001	0.539	0.557	0.514	0.402	0.199	-0.063	-0.361
Performing Operations	-0.003	-0.001	0.117	0.210	0.283	0.301	0.191	0.017	-0.201
Chemistry; Metallurgy	-0.001	0.001	0.239	0.231	0.211	0.163	0.049	-0.128	-0.349
Textiles; Paper	-0.002	-0.001	0.458	0.451	0.462	0.417	0.371	0.270	0.189
Fixed Construction	-0.002	0	0.211	0.331	0.397	0.386	0.297	0.215	0.161
Mechanical Engineering	-0.001	0	-0.456	-0.505	-0.467	-0.358	-0.208	-0.061	0.059
Physics	-0.003	-0.001	-0.156	0.016	0.212	0.343	0.284	0.068	-0.207
Electricity	-0.003	-0.002	0.373	0.465	0.540	0.580	0.552	0.367	0.070

Notes: This table shows the trends of the technology waves, HHI, New-to-Incumbent ratio and the detrended cross correlations among them. The trend is derived by running linear regressions of the focal time series on year and taking the coefficient; the cross correlations are derived by computing the correlation coefficients at different year gaps of the detrended time series.

B.7 First-Stage Results and Instrument Relevance

Table 12 reports the first-stage results for the 2SLS regressions. Columns (1)–(2) relate the contemporaneous Tech Wave Index to its shift-share instrument in the HHI regression at the 6-digit NAICS-by-year level. Columns (3)–(4) and (5)–(6) report the corresponding first-stage relationships for the one-year and two-year lagged Tech Wave Index, respectively. Columns (7)–(8) show the first-stage relationship between the contemporaneous Tech Wave Index and its instrument in the patent-level regression with the new-firm indicator as the outcome. The last row reports the first-stage F-statistics assessing instrument relevance. Across all specifications, these statistics indicate a strong

and statistically significant relationship between the instrument and the endogenous regressor.

Table 12: First-stage Results of 2SLS Regressions

	HHI						New Firm Indicator	
	$TWI_{s,t}$		$TWI_{s,t-1}$		$TWI_{s,t-1}$		$TWI_{c,t}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\lambda_{s,t}$	0.571***	0.736***						
	(0.157)	(0.0874)						
$\lambda_{s,t-1}$			0.345*	0.719***				
			(0.183)	(0.0849)				
$\lambda_{s,t-2}$					0.347*	0.722***		
					(0.184)	(0.0864)		
$\lambda_{c,t}$							0.600***	0.707***
							(0.121)	(0.0999)
Observations	37500	37500	37500	37500	36500	36500	4032000	4032000
Industry/IPC FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	No	Yes	No	Yes	No	Yes	No	Yes
F excl. instruments	13.21	70.85	3.54	71.86	3.56	69.92	24.53	49.99

Notes: Standard errors are clustered at the 6-digit NAICS-by-year level for columns (1)–(6), and at the 4-digit IPC-by-year level for columns (7)–(8). The inclusion of year and industry fixed effects, and the F-statistics assessing instrument relevance are shown in the last three rows. To comply with Census Bureau disclosure requirements, the number of observations is rounded to the nearest hundred. *** Significant at the 1 percent level; ** significant at the 5 percent level; * significant at the 10 percent level.

B.8 Alternative Measures of Market Concentration

The main text of this paper uses the Herfindahl-Hirschman Index to measure market concentration. It captures the whole distribution of firm sales in the economy, but the limitation is that it is based on only publicly listed firm. An alternative measure of market concentration is the share of sales by the top firms. This paper adopts the cleaned data series by Kwon, Ma and Zimmermann (2023) to calculate respectively the three-year moving average of the receipt share of the top 0.1% and 1% firms. The top shares are generated by the IRS data, which covers a more comprehensive set of firms. So, it can be used as a complement to the HHI measure in the paper. As displayed in Figure 16, the top shares exhibit increasing trends in general but with fluctuations. The peaks and troughs of the fluctuations appear nearly simultaneously with the HHI measured in this paper, showing the robustness of the market concentration patterns shown in the paper.

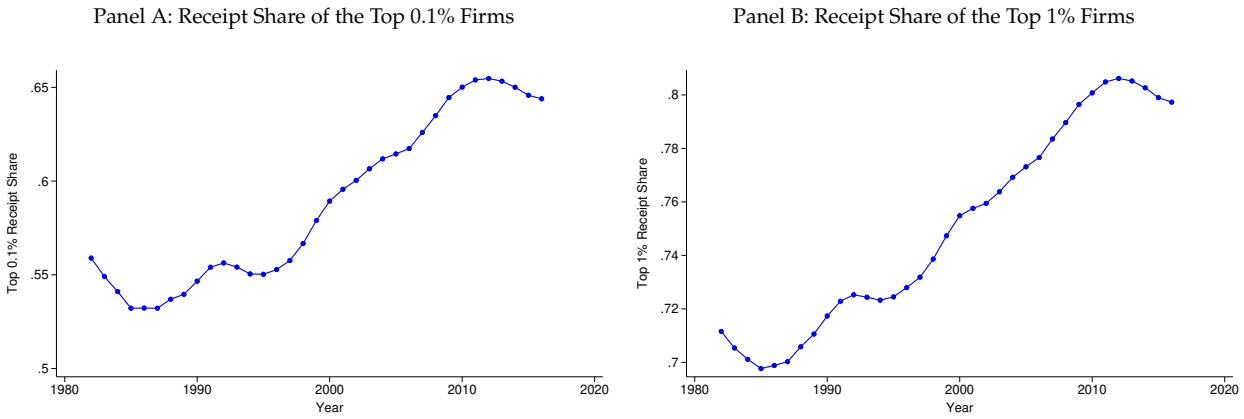


Figure 16: Receipt Shares by Top firms

Notes: This figure shows the three-year moving average of the receipt share of the top 0.1% (Panel A) and 1% firms (Panel B). The receipt shares are from the cleaned data series by Kwon, Ma and Zimmermann (2023), which is posted on https://businessconcentration.com/. The data source is the Statistics of Income (SOI) and the associated Corporation Source Book published annually by the IRS. Their statistics cover the whole population of US corporations.
Sources: https://businessconcentration.com/.

B.9 IPC-Level Relationship between Tech Waves and New-to-Incumbent Ratio

To assess the robustness of the relationship between idea allocation and technology waves, this paper compares the two trends by patent technological fields, categorized by the first digit of the patent IPC code. The IPC-level "Novelty" Index and "New-to-Incumbent Ratio" are computed using the same methodology as described in equations 1 and 4, with patent sets segregated according to their respective technology classes. Figure 17 illustrates that a positive correlation between idea allocation and technology waves is consistently observed across most technology classes. When a specific technology class experiences breakthroughs, there is an increase in the flow of ideas toward new startups. The contemporaneous correlation coefficients between the two curves are, respectively, 0.40 for Human Necessities, 0.30 for Performing Operations, 0.16 for Chemistry, 0.42 for Textiles, 0.39 for Fixed Constructions, -0.47 for Mechanical Engineering, 0.34 for Physics, and 0.58 for Electricity. The cross correlations when $k \in \{-3, -2, -1, 0, 1, 2, 3\}$ for each technological field are shown in Table 11 in Appendix B.6.

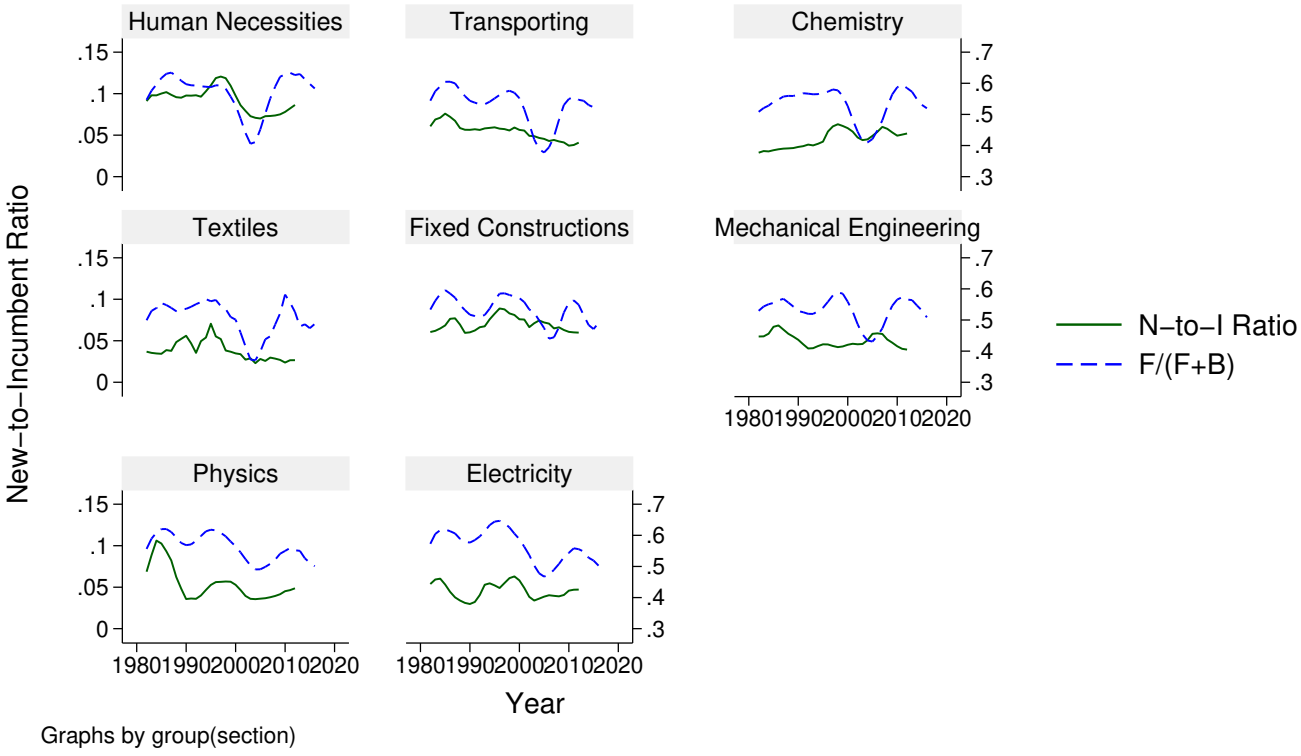


Figure 17: technology waves and Idea Allocation by Patent Technology Class

Notes: This figure shows the technology waves and the idea allocation between new and incumbent firms by patent technology class. The blue dashed curve, based on the methodology defined in this paper, calculates the relative ratio of forward citations to the sum of forward and backward citations. The green solid curve displays the "New-to-Incumbent Ratio" defined in the paper, capture where new ideas contribute their value. The two curves have different y-axises, shown respectively on the left and right.
Sources: Longitudinal Business Database (LBD) and USPTO patent and citation data.

C Model and Proof

C.1 Starting up New Businesses

The partner's problem has the same form as the incumbent firm's, with $\tilde{q} = 0$ and the average innovation value being z_0 instead of $x_0(z_0, \tilde{q})$ (Equation (17)).

$$\begin{aligned} \max_a (1-a)\lambda_0 e_I \int V(j,t|zQ(t))dz - T_I V(t|Q(t) \\ \text{st } e_I = \arg\max \{u(c_I(a,0,T_I), e_I)\} \\ u(c_I(a,T_I,0,Q(t)), e_I) \geq \bar{u}(z_0) \\ (1-a)\lambda_0 e_I \int V(j,t|zQ(t))dz - T_I V(t|Q(t) \geq 0 \end{aligned} \tag{42}$$

The partners are assumed to get zero net payoff due to competition.

The inventor decides her effort level by maximizing her utility, which yields:

$$e_I = \lambda_0 a z_0 - A a^2 \lambda_0 \mathbb{E}\left(z(z_0)^2\right) \tag{43}$$

The firm's problem in Equation 42 becomes

$$\begin{aligned} \max_a \lambda_0 e_I z_0 - A a^2 \left(\lambda_0 e_I \mathbb{E}\left(z(z_0)^2\right)\right) - \frac{1}{2} e_I^2 \\ \text{st } e_I = \lambda_0 a z_0 - A a^2 \lambda_0 \mathbb{E}\left(\tilde{z}(z_0)^2\right) \end{aligned} \tag{44}$$

It gives the highest utility an inventor can obtain when working in a startup.

C.2 Production: Profit Maximization

The production sector features two types of firms: a representative final goods producer and intermediate goods producers. The final good producer assembles intermediate goods, denoted by j within the range $[0, N_F(t)]$, to produce final goods. The final goods producer chooses $\{y(t)_j\}_j$ to maximize its profit using the technology described in Section 3.2. The subscript t is omitted in this section whenever it does not cause a confusion. The final goods producer's problem can be written as:

$$\max_{\{y_j\}} \frac{1}{1-\beta} \int_0^{N_F} q_j^\beta y_j^{1-\beta} dj - \int_0^{N_F} y_j p_j dj. \tag{45}$$

The first-order condition

$$p_j = q_j^\beta y_j^{-\beta}$$

yields the demand function for goods produced by intermediate firms.

The intermediate goods are produced by their corresponding firm $j \in [0, N_F]$ using only labor $y_j = Ql_j$. Intermediate good producers engage in monopolistic competition, optimizing their profit by choosing l_j, p_j, y_j, given the wage level w:

$$\begin{aligned} &\max_{l_j, p_j, y_j} y_j p_j - w l_j. \\ &s.t.\ y_j = Q l_j \\ &\quad p_j = q_j^{\beta} y_j^{-\beta} \end{aligned} \tag{46}$$

The FOC yields:

$$y_j(t) = q_j(t) \left(\frac{Q(t)\,(1-\beta)}{w(t)} \right)^{\frac{1}{\beta}}, l_j(t) = \frac{y_j(t)}{Q(t)}, p_j(t) = \frac{w(t)}{Q(t)\,(1-\beta)}.$$

The labor market clears, which derives that

$$\frac{\int_0^{N_F} q_j \left(\frac{Q(1-\beta)}{w} \right)^{\frac{1}{\beta}} dj}{Q} = 1$$

Therefore, the wage is $w = (1-\beta)Q$.

C.3 Firm-led Innovation New

For firm-led innovation, plug the linear value function (Equation 33) into the firm choice problem specified in Equation 22

$$\begin{aligned} \Omega_F(z_0, q, t) &= \max_{e_F, T_F} \lambda_0 e_F \int \kappa x_F(z_0) \nu(t) N_F(t) q dx - \frac{1}{1+\delta} e_F^{\delta+1} \nu(t) N_F(t) q - T_F \nu(t) N_F(t) q \\ &= \max_{e_F, T_F} \left(\lambda_0 e_F \kappa x_{F0} - \frac{1}{1+\delta} e_F^{\delta+1} - T_F \right) \nu(t) N_F(t) q \end{aligned} \tag{47}$$

FOC yields:

$$e_F = (\lambda_0 \kappa x_{F0})^{\frac{1}{\delta}},$$

which is invariant to time and firm size. Equation 23 can be written as:

$$\lambda_0 e_F \kappa x_{F0} \nu(t) - \frac{1}{1+\delta} e_F^{\delta+1} \nu(t) = T_F \nu(t)$$

which yields:

$$\begin{aligned} T_F =& \lambda_0 e_F \kappa x_{F0} - \frac{1}{1+\delta} e_F^{\delta+1} \\ =& \frac{\delta}{1+\delta} (\lambda_0 \kappa x_{F0})^{\frac{1+\delta}{\delta}} \end{aligned} \tag{48}$$

which is also invariant to time and firm size. Inventor's utility in this case is:

$$U_I(T_F(z_0, q), 0) = T_F V(t|N_F(t)q) = \frac{\delta}{1+\delta} (\lambda_0 \kappa x_{F0})^{\frac{1+\delta}{\delta}} \nu(t) N_F(t) q,$$

which is linear in $\nu(t)N_F(t)$.

C.4 Inventor-led Innovation

In the inventor-led case, the consumption is

$$\begin{aligned} & c_I(a, T_I, q, z_0, Q(t))dt \\ & = T_I V(t|Q(t))dt + adV(j, t|q, z_0), \\ & = T_I \nu(t) Q(t) dt \begin{cases} +a\tau\nu(t)qdt, & pr = 1 - \tau dt - \lambda_0 e_I dt \\ -a(1 - \tau dt - rdt)\nu(t)q, & pr = \tau dt \\ +a((1 - rdt)x(z_0, \tilde{q})Q(t) + \tau q dt)\nu(t), & pr = \lambda_0 e_I dt \end{cases} \end{aligned} \tag{49}$$

which is linear in $\nu(t)$. The expected consumption is

$$\mathbb{E}\left(c_I(a, T_I, q, z_0, Q(t))\right) = (a\lambda_0 e_I x_0\left(z_0, \tilde{q}\right) + T_I)\nu(t)Q(t),$$

and the associated variance is

$$\begin{aligned} \text{Var}\left(c_I\left(a, T_I, \tilde{q}, z_0, Q(t)\right)dt\right) =& \mathbb{E}((c_I dt)^2) - \mathbb{E}(c_I dt)^2 \\ =& \mathbb{E}\left(a^2\nu(t)^2 Q(t)^2\left(\tilde{q}^2\tau dt + x^2(z_0, \tilde{q})\lambda_0 e_I dt\right)\right) \\ & - \mathbb{E}\left(a\nu x_0(z_0, \tilde{q})Q(t)\lambda_0 e_I dt\right)^2 \\ =& a^2\nu(t)^2 Q(t)^2\left(\tilde{q}^2\tau + \mathbb{E}(x^2(z_0, \tilde{q}))\lambda_0 e_I\right)dt \end{aligned}$$

The variance comes from two sources: non-innovation-related firm value and the R&D process. Both terms increases in firm size $\tilde{q}$, but the former one increases in a faster speed, implying that in larger firms, shocks unrelated to R&D are stronger. Hence, larger firms are subject to larger incentive problems and the equity held by the inventor provides a

weaker incentive for R&D efforts. Upon reviewing all available contracts, an inventor determines her preferred firm $\tilde{q}$.

The instant utility is:

$$U_I(c_I(t), e_I(t)) = \left((a\lambda_0 e_I x_0 (z_0, \tilde{q}) + T_I) - Aa^2 \left(\tilde{q}^2 \tau + \mathbb{E}(x^2(z_0, \tilde{q}))\lambda_0 e_I \right) - R(e_I) \right) \nu(t) Q(t) \tag{50}$$

The inventor's problem is to maximize her utility:

$$\max_{e_I(t)} U_I(c_I(a, T_I, q, Q(t)), e_I(t)).$$

FOC yields:

$$e_I(t) = \left(a\lambda_0 x_0 (z_0, \tilde{q}) - Aa^2 \mathbb{E}(x^2(z_0, \tilde{q}))\lambda_0 \right)^{\frac{1}{\delta}} \tag{51}$$

which is invariant to $\nu(t)$.

The wage rate T_I is determined by the Betrand competition:

$$(1-a)\lambda_0 e_I \left[\int V(j,t|q + x(z_0, \tilde{q}(t))Q(t))dx - V(j,t|q) \right] - T_I V(t|Q(t)) = 0$$
$$T_I = (1-a)\lambda_0 e_I x_0(z_0, \tilde{q}(t))$$

which is invariant to $\nu(t)$. Therefore, the instant utility is linear in $\nu(t)Q(t)$:

$$\begin{aligned} U_I(c_I(t), e_I(t)) &= \left((a\lambda_0 e_I x_0 (z_0, \tilde{q}) + T_I) - Aa^2 \left(\tilde{q}^2 \tau + \mathbb{E}(x^2(z_0, \tilde{q}))\lambda_0 e_I \right) - R(e_I) \right) \nu(t) Q(t) \\ &= \left(\lambda_0 \left(\left(1 - \frac{a}{\delta + 1} \right) x_0 (z_0, \tilde{q}) - \frac{\delta}{\delta + 1} Aa^2 \mathbb{E}(x^2(z_0, \tilde{q})) \right) e_I - Aa^2 \tilde{q}^2 \tau \right) \nu(t) Q(t) \\ &= u_I(z_0, \tilde{q}) \nu(t) Q(t) \end{aligned} \tag{52}$$

where $u_I(z_0, \tilde{q}) = \left(\lambda_0 \left(\left(1 - \frac{a}{\delta+1}\right) x_0 (z_0, \tilde{q}) - \frac{\delta}{\delta+1} Aa^2 \mathbb{E}(x^2(z_0, \tilde{q})) \right) e_I - Aa^2 \tilde{q}^2 \tau \right)$, invariant to $\nu(t)$.

The inventor's choice of firm, as specified in Equation (24), can be written as:

$$u(z_0, t) = \max\{ u_I(z_0, \tilde{q}^*) \nu(t) Q(t), u_I(z_0, 0) \nu(t) Q(t) \} \tag{53}$$

Since $\nu(t)Q(t) > 0$, the decision rule is the same as:

$$u(z_0, t) = \max\{ u_I(z_0, \tilde{q}^*), u_I(z_0, 0) \}. \tag{54}$$

Thus, the decision rule $\tilde{q}^*(z_0)$ does not change over time.

C.5 Proof of Proposition 1 and Lemma 1

Proof. Production is linear in aggregate quality $Q(t)$:

$$Y(t) = \frac{Q(t)}{1-\beta} \quad \Rightarrow \quad \frac{\dot{Y}(t)}{Y(t)} = \frac{\dot{Q}(t)}{Q(t)} \equiv g.$$

Since $C_I(t)/Y(t)$ and $R_F(t)/Y(t)$ are linear functions of $\nu(t)$ and $Y(t) \propto Q(t)$, define

$$s_H(t) \equiv \frac{C_H(t)}{Y(t)}, \qquad s_I\nu(t) \equiv \frac{C_I(t)}{Y(t)}, \qquad s_F\nu(t) \equiv \frac{R_F(t)}{Y(t)},$$

so the resource constraint gives $s_H(t) = 1 - (s_I + s_F)\nu(t) \equiv S_H(\nu(t))$. Household optimization implies the Euler equation

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

Because $C_H(t) = Y(t)s_H(t)$, we have

$$\frac{\dot{C}_H(t)}{C_H(t)} = \frac{\dot{Y}(t)}{Y(t)} + \frac{\dot{s}_H(t)}{s_H(t)} \quad \Rightarrow \quad r(t) = \rho + g + \frac{d}{dt}\ln s_H(t). \tag{55}$$

Plugging $V(j,t \mid q) = \nu(t)\, q$ into the HJB gives

$$\dot{\nu}(t) = (r(t) + \tau)\nu(t) - \beta.$$

Substitute $r(t)$ from (55) to obtain

$$\dot{\nu}(t)\Big(1 - \nu(t)\frac{S_H'(\nu(t))}{S_H(\nu(t))}\Big) = \big(\rho + \tau + g\big)\nu(t) - \beta. \tag{56}$$

Define

$$D(\nu) \equiv 1 - \nu\,\frac{S_H'(\nu)}{S_H(\nu)}, \qquad \Phi(\nu) \equiv \big(\rho + \tau + g\big)\nu - \beta. \tag{57}$$

According to (35) and (36), $s_I \geq 0$ and $s_F \geq 0$. Hence,

$$S_H'(\nu) = -(s_I + s_F) \leq 0, \qquad S_H(\nu) \in (0,1) \quad \Rightarrow \quad D(\nu) \geq 1.$$

Equation (56) is therefore the autonomous scalar ODE

$$\dot{\nu}(t) = \frac{\Phi(\nu(t))}{D(\nu(t))}. \tag{58}$$

Because $g > 0$ and $\rho + \tau > 0$, so $\Phi'(\nu) = \rho + \tau + g > 0$. Since $\Phi(0) = -\beta < 0$, $\lim_{\nu\to\infty} \Phi(\nu) = +\infty$, hence there exists a unique $\nu^\star > 0$ with $\Phi(\nu^\star) = 0$.

Next, solving $\dot{\nu} - (r + \tau)\nu = -\beta$ by an integrating factor and imposing the transversality condition derive the present-value representation

$$\nu(t) = \int_t^\infty \beta \, \exp\Big(- \int_t^s \big(r(u) + \tau\big)\, du\Big)\, ds.$$

Since $r(u) = \rho + \dot{C}_H/C_H \geq \rho$, we have $r(u) + \tau \geq \rho + \tau$ and thus

$$0 < \nu(t) \leq \int_t^\infty \beta\, e^{-(\rho+\tau)(s-t)}\, ds = \frac{\beta}{\rho + \tau} \quad \text{for all } t. \tag{59}$$

From (58), $\text{sign}(\dot{\nu}) = \text{sign}(\Phi(\nu))$. If $\nu(t_0) > \nu^\star$ at some t_0, then $\Phi(\nu(t)) > 0$ thereafter and $\nu(\cdot)$ is strictly increasing, which contradicts the upper bound in (59). If $\nu(t_0) < \nu^\star$, then $\Phi(\nu(t)) < 0$ thereafter and $\nu(\cdot)$ is strictly decreasing and bounded below by 0, so it must converge to a limit $L \in [0, \nu(t_0))$ with $\Phi(L) \leq \Phi(\nu(t_0)) < 0$, which implies $\dot{\nu} < 0$ near L, a contradiction. Hence every admissible equilibrium path satisfies

$$\nu(t) \equiv \nu^\star \qquad \text{and} \qquad g(t) \equiv g^\star.$$

With ν constant, e_F and the inventor-side policies are time-invariant in Q-normalized units, so s_I and s_F are constant and $s_H = 1 - (s_I - s_F)\nu^*$ is constant. $\dot{s}_H = 0$ in (55) implies

$$r(t) = \rho + g^\star.$$

Because $C_H(t) = Y(t)s_H(t)$, $C_I(t) = Y(t)s_I(t)$, $R_F(t) = Y(t)s_F(t)$, and $Y(t) \propto Q(t)$, all aggregates grow at the common constant rate $g^\star$. The fixed point $\nu^\star$ is unique, so the competitive equilibrium is unique and coincides with the balanced-growth path. Besides, there is no transition dynamics. □

C.6 Closed-Form Model Solution

The firm's problem in Equation 17 can be rewritten as:

$$\begin{aligned} \max_{a} & \left(\lambda_0 e_I x_0(z_0,\tilde{q}) - Aa^2\left(\tau\tilde{q}^2 + \lambda_0 x_0\left(z_0,\tilde{q}\right)^2\right) - \frac{1}{2}e_I^2\right)\nu Q \\ \text{st}\, & e_I = \lambda_0 a x_0\left(z_0,\tilde{q}\right) \end{aligned} \tag{60}$$

Putting the expression of e_I into the maximization problem and taking the FOC with regard to the equity share, a, derives,

$$\begin{aligned} a^* &= \frac{\lambda_0^2 x_0\left(z_0,\tilde{q}\right)^2\nu^2}{\lambda_0^2 x_0\left(z_0,\tilde{q}\right)^2\nu^2 + 2A\left(\tau\tilde{q}^2\nu^2 + \lambda_0 x_0\left(z_0,\tilde{q}\right)^2\nu^2\right)} \\ &= \frac{1}{1+2\frac{A}{\lambda_0}\left(\frac{\tau\tilde{q}^2}{\lambda_0 x_0(z_0,\tilde{q})^2}+1\right)} \end{aligned} \tag{61}$$

Upon reviewing all contracts, an inventor with idea quality z_0 chooses which firm $\tilde{q}$ to work for by maximizing her instant utility:

$$\begin{aligned} \max_{\tilde{q}} u\left(c_I,e_I\right) =& \mathbb{E}\left(c_I\right) - \frac{A}{\nu Q(t)}\text{Var}\left(c_I\right) - R\left(e_I\right)\nu Q(t) \\ \text{st}\, & a = a^*\left(\tilde{q}\right) \\ & T_I = T_I^*\left(\tilde{q}\right) \end{aligned} \tag{62}$$

Putting the expression of the optimal equity level, $a^*(\tilde{q})$, and $T_I^*(\tilde{q}) = (1-a^*)\,\lambda_0 e_I x_0\left(z_0,\tilde{q}\right)$ into the maximization problem and solving the first-order condition,

$$\frac{\partial x_0\left(z_0,\tilde{q}\right)}{\partial\tilde{q}} = \frac{2A\tau}{4A\tau\frac{\tilde{q}}{x_0(z_0,\tilde{q})} + \frac{2A\lambda_0+\lambda_0^2}{\tilde{q}/x_0(z_0,\tilde{q})}}, \tag{63}$$

derive the optimal firm size,

$$\tilde{q}^* = \left(\frac{\left(2A\lambda_0+\lambda_0^2\right)\left(\gamma\left(z_0\right)\right)^2 b}{2A\tau q_0^{2b}\left(1-2b\right)}\right)^{\frac{1}{2-2b}}.$$

The left-hand side and right-hand side of the first-order condition when $b < 0.5$ and $\eta = -1$ are shown in Figure 18.

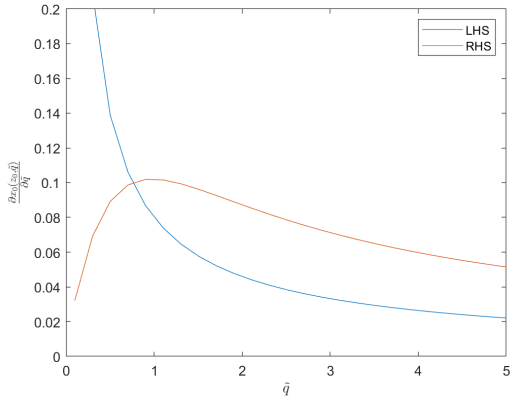


Figure 18: FOC Condition

Notes: This figure shows respectively the left-hand side (the blue curve) and the right-hand side (the red curve) of Equation 63. When $b < 0.5$, there exists a unique intersection.

C.7 Proof for Proposition 2

Proof.

$$\begin{aligned}\frac{\partial \tilde{q}^*}{\partial z_0} &= \frac{\tilde{q}^*}{\gamma(z_0)(1-b)} \frac{\partial(\gamma(z_0))}{\partial z_0} \\ &= \frac{\tilde{q}^*}{\gamma(z_0)(1-b)} \left(z_0^{\eta} + B^{\eta}\right)^{\frac{1}{\eta}-1} z_0^{\eta-1}\end{aligned}$$

where

$$\tilde{q}^* = \left(\frac{\left(2A\lambda_0 + \lambda_0^2\right)(\gamma(z_0))^2 b}{2A\tau q_0^{2b}(1-2b)}\right)^{\frac{1}{2-2b}}.$$

When $b < 0.5$, the optimal size increases in the inventor's idea quality, z_0. □

C.8 Proof for Proposition 3

Proof. The highest expected utility $u_N(z_0)$ an inventor z_0 can obtain when working in a new firm is

$$\begin{aligned}u_N(z_0) =& u\left(c_I(z_0, \tilde{q}), e_I(z_0, \tilde{q})\right) \\ =& \frac{1}{2}\frac{\lambda_0}{\lambda_0 + 2A}\lambda_0^2 z_0^2 \nu Q,\end{aligned}$$

which is quadratic in the idea quality $z0$. However, the highest expected utility $u_I(z_0)$ an inventor z_0 can obtain when working in an incumbent firm depends on B:

$$\begin{aligned} u_I(z_0) =& u\left(c_I(z_0,\tilde{q}^*), e_I(z_0,\tilde{q}^*)\right) \\ =& \frac{1}{2}\lambda_0^2 \frac{\left(\frac{2A\lambda_0+\lambda_0^2}{2(1-2b)A\tau}b\right)^{\frac{b}{1-b}}}{\left(1+\frac{2A(1-b)+\lambda_0 b}{\lambda_0(1-2b)}\right)q_0^{\frac{2b}{1-b}}}\gamma(z_0)^{\frac{2}{1-b}}\nu Q \\ =& \frac{1}{2}\lambda_0^2\hat{q}_0\gamma(z_0)^{\frac{2}{1-b}}\nu Q, \end{aligned}$$

where $\hat{q}_0$ is a parameter ($\hat{q}_0 = \frac{\left(\frac{2A\lambda_0+\lambda_0^2}{2(1-2b)A\tau}b\right)^{\frac{b}{1-b}}}{\left(1+\frac{2A(1-b)+\lambda_0 b}{\lambda_0(1-2b)}\right)q_0^{\frac{2b}{1-b}}}$). The derivative is proportional to $\gamma(z_0)^{\frac{1+b}{1-b}}\left(\frac{\gamma(z_0)}{z_0}\right)^{1-\eta}$. When $\eta < 0$, the change of $u_I(z_0)$ in z_0 is much smaller than the change of $u_N(z_0)$.

An inventor decides whether to join a startup by comparing the incumbent-startup utility $u_I(z_0) - u_N(z_0)$ and zero. The utility gap decreases in z_0, meaning that a larger share of inventors would choose incumbent firms when the entire distribution of z_0 shifts to the right. □

C.9 Proof of Proposition 4

Proof. An inventor decides whether to launch a startup by comparing the highest utility offered by incumbents $u_I(z_0)$ and startups $u_N(z_0)$. When $\frac{\gamma(z_0)^{\frac{1}{1-b}}}{z_0} < \left(\frac{\lambda_0}{\hat{q}_0(\lambda_0+2A)}\right)^{\frac{1}{2}}$, $u_I(z_0) < u_N(z_0)$, inventor chooses to initiate a startup. When $\eta < 0$, if $b < \frac{\min\left(z_0^{-\eta}\right)}{\min\left(z_0^{-\eta}\right)+\max(B^{-\eta})}$, $b - \frac{B^\eta}{z_0^\eta+B^\eta} < 0$ always holds:

$$\frac{\partial\left(\gamma(z_0)^{\frac{1}{1-b}}z_0^{-1}\right)}{\partial z_0} = \frac{\gamma(z_0)^{\frac{1}{1-b}}}{z_0^2}\frac{1}{1-b}\left(b - \frac{B^\eta}{z_0^\eta+B^\eta}\right) < 0,$$

since $\gamma(z_0) = \left(z_0^\eta + B^\eta\right)^{\frac{1}{\eta}}$. $\gamma(z_0)^{\frac{1}{1-b}}z_0^{-1}$ monotonically decreases in z_0, when holding B constant. It implies there exists a cutoff $\bar{z}_0(B)$, when $z_0 > \bar{z}_0(B)$,

$$\frac{\gamma(z_0)^{\frac{1}{1-b}}}{z_0} < \left(\frac{\lambda_0}{\hat{q}_0(\lambda_0+2A)}\right)^{\frac{1}{2}}$$

always holds, and hence $u_I(z_0) < u_N(z_0)$, inventors opts for new businesses instead of incumbent firms. □

C.10 Full Model

The firm's problem in Equation 17 becomes

$$\begin{aligned} \max_a & \left(\lambda_0 e_I x_0(z_0, \tilde{q}) - Aa^2 \left(\tau \tilde{q}^2 + \lambda_0 e_I \mathbb{E}\left(x(z_0, \tilde{q})^2 \right) \right) - \frac{1}{2} e_I^2 \right) \nu Q \\ & \text{st } e_I = \lambda_0 a x_0(z_0, \tilde{q}) - Aa^2 \lambda_0 \mathbb{E}\left(x(z_0, \tilde{q})^2 \right) \end{aligned} \tag{64}$$

FOC yields the optimal a solves a cubic equation [34]:

$$FOC = 2A^2\lambda_0^2 E(x^2)^2 a^3 - \left(\lambda_0^2 E(x)^2 + 2A\lambda_0^2 E(x)E(x^2) + 2A\tau q^2 \right) a + \lambda_0^2 E(x)^2 = 0$$

Characterize the optimal allocation a^*. The solutions to the FOC is $a_1 \leq a_2 \leq a_3$. When $a \to -\infty, f(a) < 0$, and $a = 0, f(a) > 0$. Therefore, there is at least one solution satisfying $a_1 < 0$. Consider the SOC. It can be written in this form: $f'(a) = C_1 a^2 + C_2, C_1 = 6A^2\lambda_0^2 E(x(z_0, \tilde{q})^2)^2 > 0, C_2 = \left(\lambda_0^2 E(x(z_0, \tilde{q}))^2 + 2A\lambda_0^2 E(x(z_0, \tilde{q}))E(x(z_0, \tilde{q})^2) + 2A\tau\tilde{q}^2 \right) > 0$. Obviously, there exists a cutoff $\bar{a} > 0$, such that $SOC > 0$ when $a \in (-\infty, -\bar{a}) \bigcup (\bar{a}, \infty)$, and $SOC \leq 0$ when $a \in [-\bar{a}, \bar{a}]$ Therefore, there is one and only one solution $a_1 < 0$, and there can be either two positive solutions $a_2, a_3 > 0$ (including the case where $a_2 = a_3$), or two solutions with a imagine imaginary. We are interested in the case where $a^* \in (0, 1)$ and $SOC < 0$, which is a_2 (the smaller one between the two positive solutions), if exists (easy to show that SOC at a_3 is always positive). Consider four cases:

1. $a_2 \in [0, 1], a_3 > 1$. This is the optimal allocation. $a^* = a_2$.
2. $a_2 \in [0, 1], a_3 < 1$. a_2 is an interior solution and $a = 1$ is a boundary case that may or may not be dominated by a_2. Therefore, we need to compare the case where $a = 1$ with a_2 to determine a^*.
3. $a_2 \notin R$, meaning that $FOC > 0, \forall a > 0$. $a^* = 1$.

[34] In startup, the decision rule becomes:

$$FOC = 2A^2 E(x^2)^2 a^3 - \left(E(x)^2 + 2AE(x)E(x^2) \right) a + E(x)^2 = 0$$

$$FOC = 2A^2 k_z^2 z^2 a^3 - (1 + 2Ak_z z)\, a + 1 = 0$$

4. $a_2 > 1, FOC > 0, \forall a \in [0,1]$. $a^* = 1$.

Given the contracts, inventor chooses which firm $\tilde{q}$ to work for by maximizing her utility, given her utility level.

$$\begin{aligned}
&\max_{\tilde{q}} u\left(c_I\left(a, T_I, q, Q(t)\right), e_I\right) \\
&\quad = \left(a\lambda_0 e_I x_0\left(z_0, \tilde{q}\right) + T_I - Aa^2\left(\tilde{q}^2\tau + \mathbb{E}(x^2(z_0,\tilde{q}))\lambda_0 e_I\right) - R\left(e_I\right)\right)\nu Q(t) \\
\text{st}\, e_I =& \lambda_0 a\left(z_0,\tilde{q}\right) x_0\left(z_0,\tilde{q}\right) - Aa\left(z_0,\tilde{q}\right)^2 \lambda_0 \mathbb{E}\left(x(z_0,\tilde{q})\left(z_0,\tilde{q}\right)^2\right)
\end{aligned}$$

C.11 Extension: startup tax

The firm's problem in Equation 17 becomes

$$\begin{aligned}
&\max_a \left(\lambda_0 e_I x_0(z_0,\tilde{q}) - Aa^2\left(\tau\tilde{q}^2 + \lambda_0 e_I \mathbb{E}\left(x\left(z_0,\tilde{q}\right)^2\right)\right) - \frac{1}{2}(1+\xi)e_I^2\right)\nu Q \\
&\text{st}\, e_I = \frac{\lambda_0 a x_0\left(z_0,\tilde{q}\right) - Aa^2\lambda_0\mathbb{E}\left(x\left(z_0,\tilde{q}\right)^2\right)}{1+\xi}
\end{aligned} \tag{65}$$

FOC yields that the optimal a solves a cubic equation for startups:

$$FOC = 2A^2\lambda_0^2 E(x^2)^2 a^3 - \left(\lambda_0^2 E(x)^2 + 2A\lambda_0^2 E(x)E(x^2)\right)a + \lambda_0^2 E(x)^2 = 0$$

The same expression as before. It means that introducing a tax only affects the results by reducing inventor effort without affecting the optimal equity choice.

D Quantification Details

D.1 Distribution: Matching the Model with Data

The distributional data are from the BDS, which provides information on the number of firms and total employment in each size bin, capped at 10,000 employees per firm. We use a spliced log-normal distribution with a Pareto tail to parametrically represent the firm size distribution in the initial period. Specifically, the distribution is assumed to be log-normal when $q < \bar{q}$ and Pareto otherwise. The cutoff is set at 1,000 employees.[35] The parameters are estimated using GMM. Moments are divided into two sets:

1. cumulative density related, $g_{cdf,i}$, and
2. average employment related, $g_{emp,j}$.

The first type of moments measures the cumulative density in each bin, while the second type matches the average employment of firms within each bin. With 10 bins, this yields 19 moments to match 4 parameters.

Numerically, the estimated distribution is used as the initial distribution. The main challenge is converting from a per-employee measure to the model measure, which is expressed in terms of Q. We proceed with the following algorithm:

1. Calculate the measure of firms, N_F, in the initial period.
2. Construct a grid of employment levels covering both the log-normal and Pareto parts of the distribution.
3. Calculate the CDF from the parametric distribution.
4. Using the CDF, compute the mass of firms in each bin, ensuring that the masses sum to the total measure of firms.
5. Calculate Q based on the grids and the mass allocation.
6. Normalize the grid using Q.

The normalized grid, together with the mass from step 4, is then used as the initial distribution. When simulating, the normalization is re-set every instant. As the simulation runs, it moves forward in very small increments of time, dt. At every single one of these tiny micro-steps, the reference grid is recalculated or re-normalized. Because the normalization grid updates constantly, instead of the distribution jumping abruptly from year to year, the grid evolves fluidly from one state to the next.

[35] We test alternative cutoffs, which do not materially affect the results.

D.2 Annual Growth Rate

The following figure compares the model-implied annual growth rate, simulated from 1986 onward, with its empirical counterpart. The empirical series is constructed by applying the HP filter to GDP per capita from the Federal Reserve Economic Data (FRED) and using the resulting trend component. The two series track each other closely, except that the final upswing occurs later in the data.

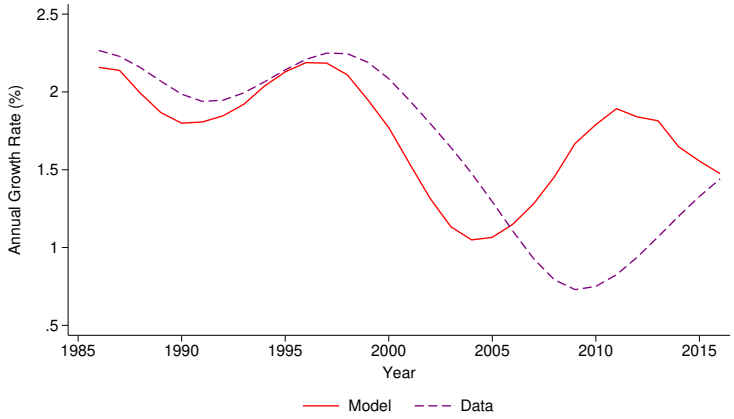


Figure 19: Annual Growth Rate

Notes: This figure reports annual economic growth rates in the model (solid line) and in the data (dashed line) from 1986 to 2016. The data series is the trend component of the annual GDP per capita growth rate, obtained using the HP filter.

D.3 Compute HHI

We compute the HHI from the simulated firm distribution defined on the discrete state grid. Our baseline measure is obtained directly from the model-implied cross-sectional distribution. Let q_{it} denote total output associated with grid point i at time t. We define $HHI_t = \sum_i (q_{it} / \sum_i q_{it})^2$. This measure corresponds to the model-implied population concentration.

As a numerical validation, we alternatively compute the HHI using a large Monte Carlo sample. In each period t, we draw firms in proportion to the total firm measure N_{Ft} and calculate the implied concentration. By the law of large numbers, the sample-based HHI converges to the population counterpart. Quantitatively, the two procedures yield indistinguishable results.

Because firms are modeled as a continuum, the level of the HHI depends mechanically on how we discretize the firm-size state space and therefore has no

standalone economic interpretation. Accordingly, we focus on trends and fluctuations, reporting the HHI normalized by its mean and emphasizing its time-series dynamics.

D.4 Robustness Checks—Different Initial Year

The following figures plot the simulated growth rates in each 5-year window, annual HHI and its extensive-margin component, and the new-to-incumbent ratio when the simulation starts in 1984. The results closely match the baseline results, which start from 1986.

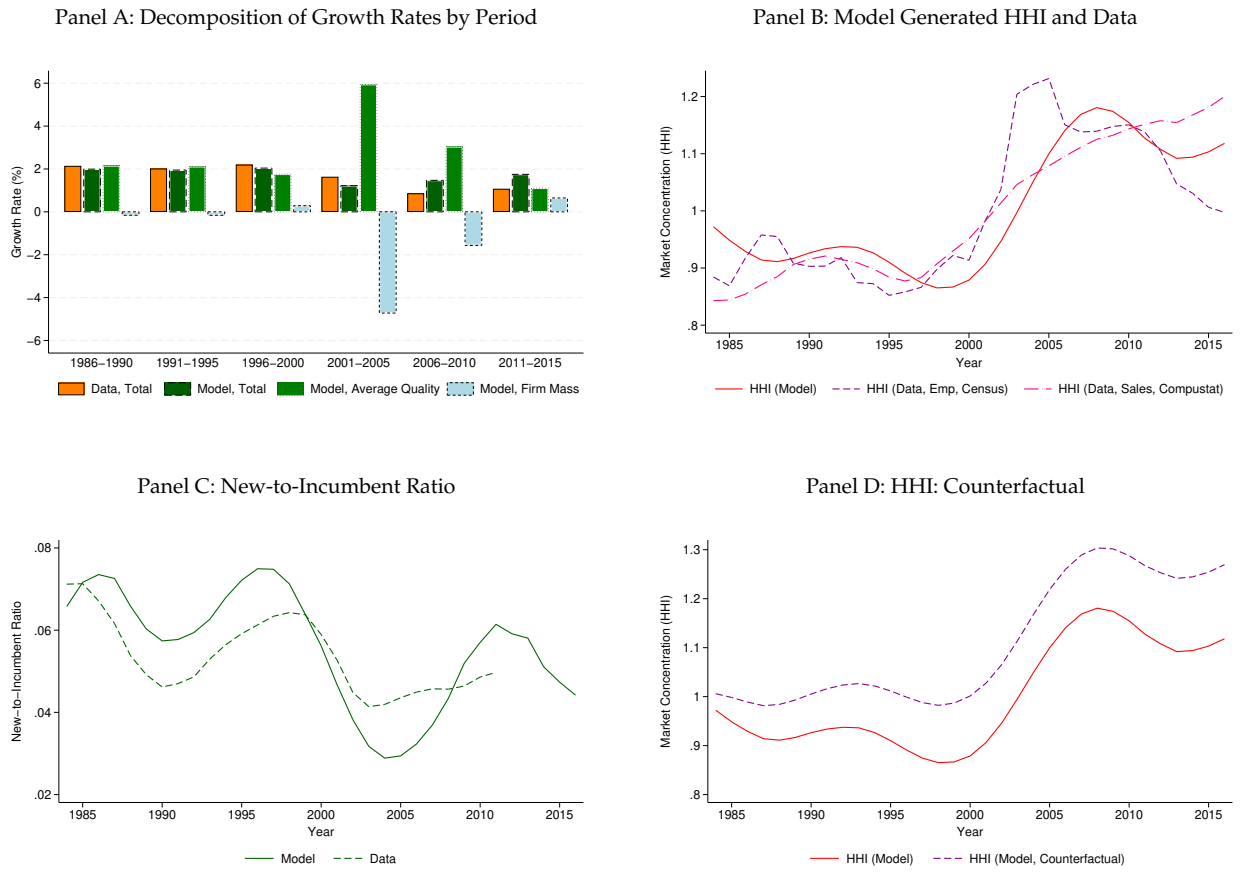


Figure 20: Simulation Results Starting from 1984

Notes: Panel A reports five-year average aggregate growth and decomposes it into improvements in firm quality and changes in firm mass. Panel B compares normalized HHI over time in the model (simulated annual HHI) and in the data (employment-based HHI from the LBD and sales-based HHI from Compustat). Panel C plots the new-to-incumbent ratio in both the model and the data. Panel D isolates the contribution of the intensive margin by comparing the full-model HHI with the counterfactual HHI. All simulations begin in 1984.

D.5 Robustness Checks—Alternative TWI Measure

To assess the robustness of our results to an alternative measure of technology waves, we re-simulate the model to match the evolution of the "RETech" index constructed by Bowen III, Frésard and Hoberg (2023). Specifically, we adjust the scale parameter of the idea-quality distribution, z_m, so that the model-implied Tech Wave Index aligns with the "RETech" series after normalizing it to have the same mean and standard deviation as our TWI measure.[36] The figures below report simulated growth rates by 5-year window and the annual HHI. The simulations continue to display a post-2000 decline in growth, driven primarily by a contraction in firm mass, and a sustained rise in market concentration consistent with the data. Overall, these results confirm that our findings are robust to using alternative measures of technology waves.

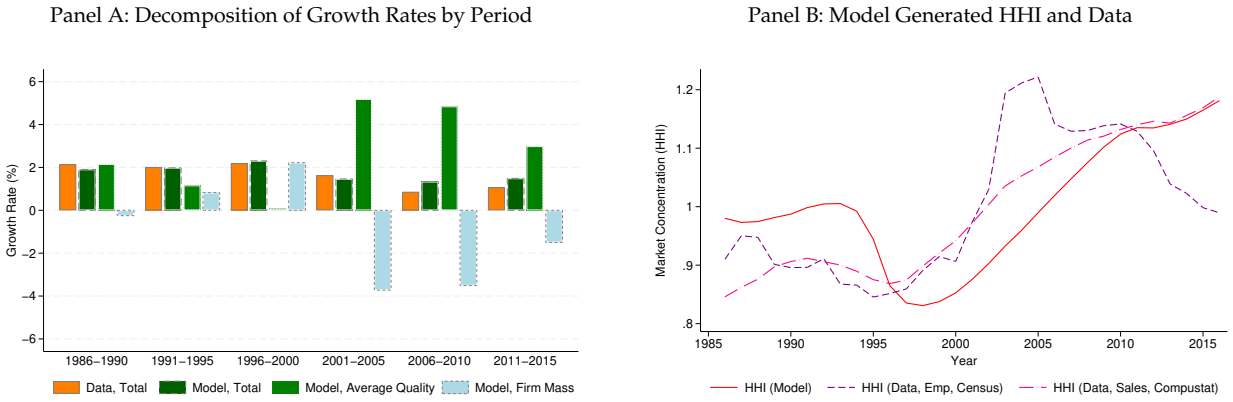


Figure 21: Simulation Results when Matching TWI with Retech Index

Notes: Panel A reports 5-year average aggregate growth and decomposes it into firm-quality improvements and changes in firm mass. Panel B compares normalized HHI over time in the model (simulated annual HHI) and in the data (employment-based HHI from the LBD and sales-based HHI from Compustat). In the simulation, we adjust the scale parameter of the idea-quality distribution, z_m, so that the model-implied Tech Wave Index matches the "RETech" index constructed by Bowen III, Frésard and Hoberg (2023).

[36] This normalization places the two technology-wave measures on a comparable scale.