Growth with Endogenous Industry Dynamics

Hyunjae Lee*

April 2026

Abstract

This paper develops a Schumpeterian growth model with endogenous industry life cycles. Potential entrants choose whether to join an existing industry or create a new one. Industries begin as monopolies, attract followers, and gradually mature. Individual industries follow non-stationary paths, yet the economy exhibits steady aggregate growth with continual sectoral rotation. The model yields a decomposition of growth into variety creation, frontier innovation, and follower catch-up, and opens a policy margin absent from conventional frameworks with an exogenous set of industries. Because entrants cannot fully appropriate the knowledge gains from new industries, the decentralized equilibrium creates too few varieties. Calibrated to U.S. data, equal-cost policy experiments suggest that subsidies to variety creation may generate larger gains in growth and welfare than subsidies to incumbent R&D in the model.

Keywords: Schumpeterian growth; Endogenous industry dynamics; Industry life-cycle; Variety creation; Innovation policy

JEL Codes: O41, O31, L16, D62, O38

*University of North Carolina at Chapel Hill. E-mail: hyunjae.lee@unc.edu. I am deeply indebted to my advisor Can Tian for her invaluable advice and guidance. Marlon Azinovic-Yang, Lutz Hendricks, Will Jungerman, Yasutaka Koike-Mori, Stan Rabinovich, Mohamed Adhami, Yongseok Shin and Roozbeh Hosseini provided valuable comments.

1 Introduction

Productivity gains are concentrated in a small set of sectors whose identity changes over time, rather than being spread uniformly across the economy (Ferguson and Wascher, 2004; Sarte and Taylor, 2025). The micro-level pattern underlying this rotation is well documented. Young industries exhibit high entry and rapid innovation, while mature industries see declining entry and more incremental innovation (Gort and Klepper, 1982; Abernathy and Utterback, 1978; Klepper, 1996; Argente, Lee and Moreira, 2024). Firm innovation and entry decisions thus depend on an industry's stage of development, and those decisions in turn shape how industries emerge and evolve. Because these innovations are fundamental drivers of aggregate productivity growth, a feedback between industry life-cycles and firm-level decisions governs the long-run growth rate. Understanding this mechanism, and the industry composition it implies, is important both for explaining why aggregate growth has remained relatively steady even as individual industries rise and decline, and for evaluating innovation policies whose effects operate through the composition of young and mature industries.

In this paper, I develop and quantify a step-by-step Schumpeterian growth model in which the set of active industries is endogenous. Industries are created by entrants, evolve as followers join and competition intensifies, and eventually exit due to obsolescence. Within each industry, the intensity of competition is itself endogenous, determined by the number of rivals and the size of technology gaps. This intensity shapes both incumbent innovation incentives and the attractiveness of the industry to new entrants. A potential entrant can either join an existing industry or pay a sunk cost to create a new one. The entry decision governs the rate at which new industries appear and shapes the cross-sectional distribution of industries over life-cycle states and the long-run growth rate.

Because entry shapes the cross-sectional distribution, endogenous industry life-cycles create policy margins that fixed-industry models cannot produce. The entry mechanism generates life-cycles as an equilibrium outcome. Individual industries follow non-stationary paths, yet the stationary cross-sectional distribution sustains steady aggregate growth with an explicit decomposition into variety creation, frontier innovation, and catch-up innovation. The framework therefore opens a policy margin absent in fixed-industry models, namely the allocation of entrants between new and existing industries. This margin exists because the cross-sectional distribution is itself an equilibrium object that responds to the entry decision. Private entrants undervalue industry creation relative to the constrained optimum, and equal-cost policy experiments calibrated to U.S. data illustrate the quantitative importance of this margin. Subsidies targeted at variety creation have larger effects on growth and welfare than subsidies to incumbent R&D, because they operate directly on the entry margin that the equilibrium undervalues.

The model generates industry life-cycles endogenously from this entry mechanism. New industries begin as small monopolies with a one-step lead over a competitive fringe. As they prove profitable, they attract followers, transition into oligopoly, and experience high innovation rates as leaders and followers compete in a patent race. Over time, leaders pull ahead, followers catch up less frequently, and the incentives for further entry and follower innovation decline. Mature industries become larger but less dynamic, and joining them becomes less attractive for new entrepreneurs. At that point, potential entrants are more likely to bear the cost of creating new industries instead.

Although individual industries follow non-stationary paths, the economy admits a stationary Markov-perfect equilibrium with an invariant cross-sectional distribution over life-cycle states and a constant long-run growth rate. The resulting equilibrium yields an explicit decomposition of the aggregate growth rate into three components: variety creation through the birth of new industries, frontier innovation by leaders within monopolies, and catch-up innovation by followers in oligopolies, each weighted by the stationary distribution over life-cycle states. This decomposition maps the cross-section of industries directly to aggregate growth, and clarifies how stable long-run growth can coexist with front-loaded innovation and declining business dynamism at the industry level.

The same entry decision that generates these industry dynamics also creates a wedge between private and social incentives to create new industries. The social return to a new industry includes the permanent expansion of the public knowledge stock, which private entrants cannot fully capture. This wedge is amplified by creative destruction, which causes private firms to discount future profits at a rate that exceeds the social discount rate. A partially offsetting force is the resource cost of market creation, which the planner values at a shadow price that accounts for the aggregate resource constraint. In the calibrated economy, the knowledge and creative-destruction forces quantitatively dominate the resource-cost force, so the equilibrium produces too little industry creation relative to the constrained optimum. The relative effectiveness of different innovation subsidies depends on how the cross-sectional distribution responds to changes in the entry decision. A policy that stimulates R&D in young, rapidly expanding industries may therefore have very different effects in mature, follower-heavy industries where business stealing is severe and innovation incentives are weak.

I discipline the magnitudes by calibrating the model to U.S. data from 1982–2011. The calibrated economy exhibits too little industry creation. Potential entrants are too willing to enter existing oligopolies to chase short-term business-stealing rents, and too reluctant to pay the fixed costs to open new markets. Comparative statics reveal that the relationship between startup entry and growth is U-shaped, reflecting the tension between business-stealing and variety-creation effects. Tougher product-market competition raises growth primarily by redirecting entrants away from mature industries and toward creating new ones. In equal-cost policy experiments, subsidies targeted at variety creation deliver larger gains in long-run growth and welfare than subsidies to incumbent R&D, because they operate directly on the entry margin that the equilibrium undervalues.

Related Literature This paper connects four literatures that have developed largely in parallel. The first studies within-industry Schumpeterian competition and quality-ladder innovation. The second studies growth through product creation and technological turnover. The third documents systematic life-cycle regularities in industries and product markets. The fourth studies innovation policy in endogenous-growth and industry-dynamics environments. My contribution is to endogenize industry life-cycles within a Schumpeterian growth framework and to show that the resulting life-cycle structure shapes aggregate growth and creates policy margins that fixed-industry frameworks cannot produce.

A large literature studies within-industry innovation races and market structure in Schumpeterian growth models. Aghion et al. (2001) introduce the step-by-step framework in which firms' innovation incentives depend on their position relative to rivals. Akcigit and Ates (2021, 2023) extend this framework to match ten facts on declining U.S. business dynamism and identify declining knowledge diffusion as the dominant driver. Liu, Mian and Sufi (2022) use a similar structure to show that low interest rates raise market power and reduce productivity growth. Olmstead-Rumsey (2022) studies how declining laggard innovativeness drives rising concentration and the productivity slowdown. Cavenaile, Celik and Tian (2025) develop an oligopolistic Schumpeterian model with an endogenously determined number of large firms within each industry and strategic innovation choices, with implications for markups and business dynamism. These models generate rich within-industry dynamics, including escape-competition effects, endogenous technology gaps, and leadership turnover. But the key entry decisions in all of them operate within existing markets. None makes the allocation of entrants between incumbent industries and entirely new industries an equilibrium outcome. The rate at which the economy creates new markets therefore lies outside the scope of their policy analysis.

A parallel literature studies growth through product creation and technological turnover. Romer (1990) introduced the expanding-variety growth model. Akcigit and Kerr (2018) distinguish internal innovation that improves existing products from external innovation that creates new product lines and find that entrants and small firms contribute disproportionately to high-impact external innovations. Garcia-Macia, Hsieh and Klenow (2019) decompose U.S. productivity growth and show that incumbent own-product improvements account for the largest share, a decomposition I return to in the policy analysis. In a related endogenous-growth model of technological turnover, Aghion et al. (2026) have entrants introduce new technologies and incumbents develop them incrementally, sustaining long-run growth through periodic replacement. Ribeiro (2026) is also related in spirit. He studies growth with continually emerging technology vintages and a stationary distribution of research effort across technologies of different ages. But his central margin is the allocation of R&D across new and old technologies, whereas this paper focuses on the allocation of entrepreneurial entry across incumbent and new industries. Most existing models of product creation and technological turnover do not jointly model the gradual transition from monopoly to competitive oligopoly, a front-loaded innovation pattern that declines over a life-cycle, and a stationary cross-sectional distribution over industry life-cycle states that maps directly into aggregate growth. This paper brings these elements together.

A third literature documents systematic life-cycle regularities in industries and product markets. Gort and Klepper (1982) and Agarwal (1998) document high entry early in the life-cycle and later consolidation, while Klepper (1996) emphasizes declining product innovation and stabilizing market shares as industries mature. Abernathy and Utterback (1978) show that product innovation is front-loaded and gradually gives way to process improvements. Argente, Lee and Moreira (2024) document that individual products peak quickly and then decline as newer varieties capture market share. Argente et al. (2025*b*) combine patent and product data and document patterns consistent with strategic patenting by market leaders, which may impede rival product introduction. Jovanovic and Wang (2025) develop a model of industry evolution in which diffusion shapes innovation incentives and industry growth over the life-cycle, and calibrate it to 18 U.S. industries, including automobiles and personal computers. My model generates these patterns endogenously from the interaction of step-by-step innovation and entry decisions.

A fourth literature studies innovation policy in endogenous-growth and industry-dynamics environments. Akcigit, Hanley and Stantcheva (2022) derive optimal nonlinear R&D subsidies with heterogeneous firms but do not distinguish the variety-creation margin from incumbent R&D. Atkeson and Burstein (2019) evaluate uniform innovation subsidies and find modest scope for raising productivity. Neither paper can address how subsidy design interacts with the composition of young and mature industries, because neither endogenizes industry creation. Cavenaile et al. (2024) are especially relevant because they study the policy implications of industry life-cycles in general equilibrium with endogenous innovation and oligopolistic competition following technological breakthroughs. My paper differs by embedding a step-by-step Schumpeterian patent-race structure in a model with an explicit entry margin between joining incumbent industries and founding new ones. The stationary distribution of industries across life-cycle states and the effects of different subsidy instruments are therefore jointly determined in equilibrium.

The rest of the paper is organized as follows. Section 2 reviews stylized facts on industry life-cycles and sectoral contributions to aggregate growth. Section 3 lays out the model. Section 4 characterizes the stationary Markov-perfect equilibrium. Section 5 derives the main theoretical results: a decomposition of aggregate growth, a constrained efficiency analysis, and a characterization of industry life-cycle dynamics. Section 6 presents the quantitative analysis and policy implications. Section 7 concludes.

2 Motivating Evidence

This section documents four stylized facts about industry dynamics and aggregate productivity growth. Each fact highlights a dimension of the tension between non-stationary industry evolution and stable aggregate growth, and each motivates a specific feature of the model developed in Section 3.

Stylized Fact 1: Aggregate productivity growth is stable, but the industries driving it rotate over time.

If individual industries slow down as they mature, steady aggregate growth requires that new high-growth industries continually replace them. This is the central observation motivating the endogenous industry creation margin in the model. A large body of evidence documents this rotation. The industries making the largest contributions to aggregate productivity growth have shifted repeatedly, from durable goods manufacturing to IT-producing industries and then toward IT-using services (Oliner, Sichel and Stiroh, 2007; Jorgenson, Ho and Stiroh, 2008), with continued sectoral reallocation in more recent data (Sarte and Taylor, 2025; Hobijn et al., 2025). These episodes are part of a broader pattern of recurring productivity surges associated with different technological and organizational forces (Ferguson and Wascher, 2004), accompanied by compositional shifts in production inputs, including the rising importance of intellectual property products (Foerster et al., 2025). In the model, this rotation emerges from a stationary cross-sectional distribution of industries over life-cycle states. Individual industries follow non-stationary paths, but the population shares in each state remain constant, delivering steady aggregate growth.

Stylized Fact 2: Industry-level output rises and prices fall, while individual product-level output declines.

This contrast disciplines the model's overlapping life-cycle structure. Industries can expand through entry and innovation while prices fall as competition intensifies, even as individual varieties are displaced by newer competitors. At the industry level, Agarwal (1998), building on the diffusion patterns documented by Gort and Klepper (1982), shows that industry evolution is typically associated with rising quantities and falling prices across a wide range of product markets. At the product level, Argente, Lee and Moreira (2024) show that individual product quantities peak quickly and then decline, with products frequently replaced by new varieties. Argente and Yeh (2022) document strong age dependence in product-level pricing moments, with entering products changing prices more frequently and by larger amounts early in the life-cycle. The model captures this contrast through overlapping life-cycles. Industries expand as entry adds firms and innovation raises frontier technology, while individual varieties lose market share to newer competitors both within and across industries.

Stylized Fact 3: Entry rates decline as industries mature.

As older industries become less attractive to entrants, potential entrants are pushed toward creating new markets, the key margin of the model. Gort and Klepper (1982) and Agarwal (1998) document that, across many product markets, the number of active firms rises strongly in early stages and then levels off or declines. Using a larger set of product markets with explicit entry and exit data, Agarwal and Gort (1996) show that industry entry rates peak in early growth stages and fall over time. Several mechanisms may amplify this decline. Argente et al. (2025*a*) highlight costly demand-side frictions in firm growth, while Argente et al. (2025*b*) document patterns consistent with strategic patenting by market leaders, suggesting an additional barrier to entry in mature markets. In the model, entry into a given industry declines endogenously as technology gaps widen and expected profits for new followers fall. Beyond a finite number of followers, further entry is never profitable (Lemma 4.2). This congestion effect makes the industry creation margin operative. As incumbent industries become less attractive, potential entrants face stronger incentives to bear the sunk cost of founding a new market.[1]

Stylized Fact 4: Innovation is front-loaded in the industry life-cycle.

Front-loaded innovation means that the economy's growth engine depends on a continuous supply of young, high-innovation industries. If new industry creation slows, the cross-section ages and aggregate innovation declines even if within-industry behavior is unchanged. Abernathy and Utterback (1978) argue that new industries begin in a "fluid" phase with frequent major product innovations, after which a dominant design emerges and

[1] To keep the analysis tractable, I do not model the late-stage decline in the number of firms. A standard extension with fixed operating costs could generate such a shakeout without changing the main mechanisms.

innovation shifts toward incremental process improvements. Klepper (1996) and Agarwal (1998) document similar trajectories, with patenting activity rising in the early years of a product market and declining in later stages. Jovanovic and Wang (2025) formalize these observations in a model of industry evolution in which innovation, diffusion, and intellectual property protection shape industry dynamics over the life-cycle, calibrating the model to U.S. industries including automobiles and personal computers. In the model, front-loaded innovation shows up as a declining innovation rate at the industry level. The decline is driven by widening technology gaps. As leaders pull ahead, both leaders and followers reduce their innovation effort, and leadership turnover becomes rare.

These four stylized facts point to a growth process in which aggregate regularities emerge from the interaction of many industries at different life-cycle stages, rather than from a representative sector. They also raise a question that the evidence alone cannot answer. Given that each industry eventually slows down, how does the continuous creation of new industries sustain aggregate growth, and does the decentralized economy get the rate of industry creation right? The model developed in Section 3 provides a framework for answering these questions.

3 Model

This section presents a continuous-time Schumpeterian growth model in which the set of active industries is endogenous. The model builds on step-by-step innovation frameworks (Akcigit and Ates, 2021, 2023; Liu, Mian and Sufi, 2022; Olmstead-Rumsey, 2022; Cavenaile, Celik and Tian, 2025) but departs from them in a key respect. Rather than treating the set of product lines as fixed, it allows industries to be created by entrepreneurs and to exit over time, generating *endogenous industry life-cycles.* Industries are born when entrepreneurs open new product markets, evolve through phases of monopoly and oligopolistic competition as followers enter and innovate, and eventually become obsolete. The central strategic tension arises from the entry decision that *potential entrants* face. They can enter an existing industry to compete for market share, or they can create an entirely new industry and obtain a temporary monopoly position. This choice between intensifying competition within existing markets and expanding the set of markets governs both aggregate growth and the structural evolution of the economy.

I first describe households and the final-good technology, then the structure of industries, and finally the innovation, entry, and exit processes. Where an assumption plays a specific role in supporting the stationary equilibrium or the growth decomposition derived in later sections, I note that role when the assumption is introduced.

3.1 Household

Time is continuous and indexed by $t \in [0, \infty)$. The economy is populated by a representative household that consumes, saves, and supplies labor. Production takes place in a continuum of industries indexed by $i \in [0, N_t]$, where the mass N_t changes over time as new industries are created and existing ones exit. The household has logarithmic preferences over a final consumption good and maximizes lifetime utility

$$U_t = \int_t^\infty e^{-\rho(s-t)} \log C_s \, ds, \tag{1}$$

where $\rho > 0$ is the rate of time preference and C_t denotes consumption at time t. The household owns aggregate wealth A_t, equal in equilibrium to the total market value of incumbent firms, and supplies one unit of labor inelastically. Log preferences also imply a closed-form welfare expression used in Section 5. Its budget constraint is

$$C_t + \dot{A}_t = w_t + r_t A_t, \tag{2}$$

where w_t is the wage and r_t is the interest rate.

The final consumption good, which serves as the numeraire, is produced with a Cobb-Douglas aggregator over industry outputs:

$$Y_t = N_t \exp\left(\frac{1}{N_t}\int_0^{N_t} \ln Y_{it} \, di\right),$$

where N_t is the mass of industries at time t and Y_{it} is the output of industry i. The aggregator implies that each industry i receives the same revenue flow $\bar{Y}_t \equiv Y_t/N_t$. Firm profits can therefore be written as functions only of within-industry state variables. This allows the aggregate economy to be summarized by a tractable cross-sectional distribution over industry states, characterized in Section 4.

3.2 Industries and Production

Unless noted otherwise, I omit time subscripts for brevity. Each industry i contains a finite number of firms, collected in a set $\mathcal{F}_i$, that differ in their technology levels and therefore in their productivity and equilibrium markups. If an industry has only one firm, it is a *monopoly*. If there are multiple firms, the industry is an *oligopoly* with one *leader* (the firm with the highest technology level) and a set of *followers* that share a common (lower) technology level. The assumption that all followers share a common technology level keeps the within-industry state to two dimensions, the technology gap m and the follower count n. It also ensures that follower catch-up maps directly into the evolution of the lowest technology tier used in the growth decomposition of Section 5. Firms produce differentiated varieties within the industry and compete in prices.

Within-industry demand and production. Industry i's output is a CES aggregate of the individual varieties produced by firms $j \in \mathcal{F}_i$:

$$Y_i = \left(\sum_{j \in \mathcal{F}_i} y_{ji}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \qquad \sigma > 1,$$

where y_{ji} is the quantity of a firm j's variety and σ is the elasticity of substitution across varieties within an industry. Industry demand for Y_i comes from the final-good sector.

Each firm j combines labor with its technology to produce:

$$y_{ji} = q_{ji} \, l_{ji}, \qquad q_{ji} = \lambda^{s_{ji}},$$

where l_{ji} is labor input, $s_{ji} \in \mathbb{R}_+$ is the technology level, and $\lambda > 1$ is the step size of the quality ladder. A one-step increase in s_{ji} raises firm j's productivity by a factor λ. Throughout the paper, I use *technology gap* to refer to the difference in technology levels, measured in steps on the quality ladder, between a firm and its closest competitors. Firms take wages and the demand system as given and set prices to maximize profits. For a given industry state (number of firms and their technology levels), price competition uniquely determines equilibrium markups and profits.

Competitive fringe in new industries. The economy features a public knowledge stock, denoted $s_t^K \in \mathbb{R}_+$, defined as the average of the lowest technology levels across all active industries. This object measures the widely available knowledge that new entrants can draw on when designing a new product. When a new industry is created at time t, it starts with a monopolist whose technology is one step ahead of the public knowledge, $s_{ji} = s_t^K + 1$, and with a competitive fringe of firms whose technology is fixed at s_t^K and who do not innovate. The fringe serves two roles. Economically, it captures the early exploratory phase of a new product line, in which many unsuccessful imitators can enter quickly using widely available technologies. Technically, the fringe ensures the monopolist faces a finite price elasticity and therefore a well-defined markup, as formalized below.

3.3 Innovation, Industry Creation, and Destruction

Firms invest in R&D to improve their technology. Innovation is uncertain and arrives at Poisson rates that depend on firms' R&D choices. At the same time, new potential entrants arrive and decide whether to join existing industries or create new ones. Industries exit exogenously. This subsection describes these processes and the implied channels of creative destruction and knowledge diffusion.

Innovation and R&D costs. Consider a firm that chooses an innovation intensity $x \geq 0$. To achieve this rate, the firm must incur a flow R&D cost $R(x)\bar{Y}$ in units of the final good, where

$$R(x) \; = \; \frac{\alpha}{\psi} x^{\psi}, \qquad \alpha > 0, \; \psi > 1. \tag{3}$$

The parameter ψ governs the curvature of the R&D technology. The scale parameter α controls the overall costliness of innovation.

Conditional on choosing x, innovations arrive according to a Poisson process with intensity x. Each successful innovation raises the firm's technology index s_{ji} by one step, $s_{ji} \;\to\; s_{ji}+1$, and thus, $q_{ji} \;\to\; \lambda q_{ji}$. For followers in an oligopoly, I assume that knowledge is shared instantaneously among all followers. When any follower innovates, the technology level of all followers in that industry increases by one step. Thus, from the perspective of an individual follower, the effective innovation intensity includes both its own effort and the efforts of its peers. Some incremental improvements are easily codified and diffuse among firms using similar, relatively mature technologies.[2]

Potential entrants and industry creation. In each existing industry, a potential entrant arrives at an exogenous Poisson rate $z > 0$. This arrival rate governs the flow of entry opportunities, while the equilibrium allocation of those opportunities between incumbent entry and new-industry creation is endogenous. Upon arrival, the potential entrant chooses between:

1. *Joining an existing industry*: the entrant adopts the lowest technology level among active producers in that industry (the followers' level in an oligopoly, or the monopolist's level in a monopoly), becoming either a new follower or a neck-and-neck competitor.

[2] I abstract from exogenous follower catch-up shocks and from sudden leapfrogging within incumbent industries. In many related step-by-step frameworks, such assumptions are used to keep technology gaps from becoming too persistent or to generate sufficient turnover. In this framework, they are not needed. Catch-up is endogenous and gradual, arising from followers' own R&D, while large gaps and crowded follower tiers endogenously depress follower values, innovation, and entry into mature industries. Stationarity is instead sustained by the continual creation of new industries, which replenishes the cross-section with young, high-growth industries and prevents the economy from being dominated by mature, low-innovation ones.

This reflects the idea that older technologies are easier to access and that entry into established markets typically occurs at the lower end of the quality spectrum.

2. *Creating a new industry*: the entrant starts a fresh industry as a monopolist at technology level $s^K + 1$, one step ahead of the current public knowledge stock s^K. Creating a new industry requires paying a sunk cost $\kappa\bar{Y}$ in units of the final good, where κ is drawn from a distribution G with support $[\underline{\kappa}, \overline{\kappa}]$, $\underline{\kappa} \geq 0$, before the arrival. A lower draw κ represents a more promising or more easily implementable idea for a new product line.

The potential entrant compares the expected value of joining an existing industry with that of paying $\kappa\bar{Y}$ to found a new industry that starts with a temporary monopoly position. This entry decision jointly determines the number of firms in each industry and the rate at which new industries appear. It therefore governs the balance between intensifying competition within incumbent markets and expanding the variety of industries in the economy. Because a new industry is born at $s^K + 1$, the public knowledge stock and the rate of industry creation are linked. As industries innovate and followers catch up, s^K gradually increases, and each successive new industry starts from a higher technological base. This link between variety creation and the knowledge stock is also the source of the knowledge externality analyzed in Section 5. Private entrants cannot appropriate the full productivity gain that a new industry confers on the economy.

Industry exit. Each active industry exits the economy at an exogenous Poisson rate $\delta > 0$. When an industry exits, all firms in that industry disappear and their products cease to be available in the final-good aggregator. Industry exit captures obsolescence of product lines due to shifts in tastes, technologies, or regulations that are not explicitly modeled. Together with the continual creation of new industries, exogenous exit prevents the cross-section from drifting toward arbitrarily old, low-innovation industries and helps sustain the stationary distribution characterized in Section 4.

Before turning to the dynamic equilibrium, it is useful to note that simple parameter restrictions nest two familiar benchmark environments. Setting $z = 0$ and $\delta = 0$ shuts down industry creation and destruction. The mass of industries is then fixed and the model reduces to a step-by-step quality-ladder framework in which growth is driven solely by within-industry innovation. If I instead set $\alpha \to \infty$ so that incumbent firms do not invest in R&D, while keeping $z > 0$, long-run growth is driven solely by the creation of new industries, as in expanding-variety models. The baseline environment studied below maintains both $\alpha < \infty$

and $z > 0$, so that aggregate growth reflects a combination of quality improvements within industries and the creation of new industries over time.

4 Equilibrium

This section characterizes a stationary and symmetric Markov-perfect equilibrium. I first establish the profit structure and knowledge diffusion channels implied by the model primitives. I then describe the value functions and innovation decisions of monopolists, leaders, followers, and potential entrants, and characterize the stationary cross-sectional distribution of industries over life-cycle states.

4.1 Profit Structure and Knowledge Diffusion

Equilibrium prices, quantities, and profits within an industry are functions of a small set of state variables. To summarize the state of an industry, I adopt the following notations. A superscript M refers to monopolies or monopolist-specific variables. $\hat{s}$ denotes the technology gap between a monopolist and the fringe. For example, $V_{\hat{s}}^{M}$ is the value of a monopolist with gap $\hat{s}$ over the fringe. A subscript (m, n) refers to an oligopoly in which the leader's technology gap over the followers is m and there are n followers. For example, Π_{mn} denotes the leader's profit in such an industry. A subscript $-mn$ refers to a representative follower in an (m, n)-industry.

Given the Cobb-Douglas aggregation in the final-good sector, each active industry earns the same revenue flow $\bar{Y} \equiv \frac{Y}{N}$ in equilibrium, and I use $\bar{Y}$ to scale firm-level variables. The equilibrium firm profits are proportional to industry revenue and expressed as $\Pi = \pi\bar{Y}$ where $\pi \in (0, 1)$ is a normalized profit. The industry price is defined as $P_i \equiv \frac{\sum_j p_{ji} y_{ji}}{Y_i}$ for industry i with prices p_{ji}, $\forall j \in \mathcal{F}_i$. Profits, output, and prices satisfy the following properties in equilibrium.

Lemma 4.1 (Properties of firm profits and industry output and price)**.** *Consider a monopoly industry born at time τ and currently with technology gap $\hat{s}$ over the fringe, and an oligopoly with gap m, n followers, and follower technology level s_f.*

1. *The monopolist's normalized profit $\pi_{\hat{s}}^{M}$ depends only on $\hat{s}$, and $\pi_{\hat{s}}^{M} \nearrow 1$ as $\hat{s} \to \infty$.*
2. *The oligopoly leader and follower's normalized profits, π_{mn} and π_{-mn}, depend only on m and n. As $m \to \infty$, $\pi_{mn} \nearrow 1$ and $\pi_{-mn} \searrow 0$; as $n \to \infty$, $\pi_{mn} \searrow 0$ and $\pi_{-mn} \searrow 0$.*

3. The wage-adjusted industry price P_i/w depends only on s_τ^K in the monopoly case and (m, n, s_f) in the oligopoly case. $P_i/w \searrow 0$ as $s_\tau^K \to \infty$ for a monopoly, and $m \to \infty$, $n \to \infty$, or $s_f \to \infty$ for an oligopoly.

Lemma 4.1 implies several economic mechanisms that underpin the main theoretical results. The monotone profit properties govern innovation incentives throughout the life-cycle. A monopolist gains from increasing its relative lead over the competitive fringe, since $\pi_{\hat{s}}^M$ is increasing in $\hat{s}$. In an oligopoly, the leader invests in innovation to widen the technology gap and capture a larger profit share, while followers invest to narrow the gap and reduce the leader's advantage. Once an industry enters the oligopoly stage, additional entry reduces profits for all incumbents, making mature, crowded industries progressively less attractive. The third part establishes that the wage-adjusted industry price P_i/w falls as technology gaps widen or follower counts grow, so that the forces driving maturation also reduce prices.

The profit structure determines the flow payoffs in each industry state. The remaining ingredient before characterizing dynamic decisions is how the aggregate state s^K evolves as industries innovate and new industries are created. Knowledge diffusion in the model operates through three channels that link firm-level innovation decisions to the evolution of the cross-sectional distribution and the growth of s^K over time. The first is spillovers among followers within an industry. When any follower innovates, all followers in that industry advance, so that the lowest technology tier rises with the combined innovative effort of the follower base. The second is adoption by potential entrants that join existing industries. These entrants enter at the incumbents' lowest technology level, transmitting knowledge embodied in active producers to a new firm. The third is the creation of new industries at $s^K + 1$, which directly raises the public knowledge stock by drawing on the economy-wide frontier and embedding it in a new product line. With the profit structure and diffusion channels in hand, the next subsection characterizes the dynamic decisions of each agent.

4.2 Value Functions and Innovation Decisions

The household's Euler equation,

$$\frac{\dot{C}_t}{C_t} = r_t - \rho,$$

will be used below to connect firm values to aggregate growth. On the firm side, let $V_{\hat{s}}^M$ denote the value of a monopolist with technology gap $\hat{s}$ over the fringe, and let V_{mn} and V_{-mn} denote the values of a leader and a follower in an oligopoly with technology gap m and n followers.

Because all flow payoffs in a given industry state scale with the common revenue level

$\bar{Y}$, it is convenient to express firm values in units of $\bar{Y}$ and work with normalized values $v \equiv V/\bar{Y}$. In a stationary equilibrium these normalized values are time-invariant, so I write $V = v\bar{Y}$ and work directly with v throughout this section (see Appendix B.1).

Monopolist. A monopolist's position is defined by the size of its technological lead over the competitive fringe. This lead determines both how much profit the monopolist can extract today and how much it stands to gain from pushing further ahead. Consider a monopolist with relative technology gap $\hat{s}$ over the fringe. Its normalized value $v^M_{\hat{s}}$ solves

$$(\rho + g_N)\, v^M_{\hat{s}} = \max_{x \geq 0} \Big\{ \pi^M_{\hat{s}} - R(x) + x\big(v^M_{\hat{s}+1} - v^M_{\hat{s}}\big) + z\,\theta_M\big(v_{01} - v^M_{\hat{s}}\big) - \delta\, v^M_{\hat{s}} \Big\}. \tag{4}$$

The left-hand side is the flow cost of the monopolist's value, which includes both time discounting at rate ρ and the dilution of normalized values as the mass of industries expands at rate g_N. On the right-hand side, the monopolist collects its current normalized profit $\pi^M_{\hat{s}}$ and invests in R&D at cost $R(x)$ to generate innovations at Poisson rate x. Each successful innovation advances the monopolist one step further ahead of the fringe, yielding a capital gain of $v^M_{\hat{s}+1} - v^M_{\hat{s}}$.

The term $z\theta^M(v_{01} - v^M_{\hat{s}})$ captures the threat of entry. θ_M denotes the probability that a potential entrant who meets a monopoly chooses to join rather than create a new industry. This probability is the same for all monopolists, because the new entrant always receives v_{01}, the neck-and-neck duopoly value, regardless of the current $\hat{s}$. At rate $z\theta^M$, a potential entrant chooses to join the monopoly as a follower rather than start a new industry, converting the monopoly into a duopoly and inflicting a capital loss of $v^M_{\hat{s}} - v_{01}$. Finally, the monopolist faces exogenous exit at rate δ, destroying its accumulated value.

The optimal innovation intensity is

$$x^M_{\hat{s}} = \left(\frac{\max\{v^M_{\hat{s}+1} - v^M_{\hat{s}},\, 0\}}{\alpha} \right)^{\frac{1}{\psi-1}}.$$

The monopolist's incentive to invest in R&D is governed entirely by the marginal value of an additional step, $v^M_{\hat{s}+1} - v^M_{\hat{s}}$.

Equation (4) and the optimal policy make transparent that the monopolist innovates to extend its technological lead and earn higher profits, but faces two forces that limit the incentive to do so. First, there are diminishing marginal gains from pushing an already-large lead further, as Lemma 4.1 establishes. Second, entry may convert the market into competition and wipe out the monopoly position. These forces drive the eventual slowdown of innovation as industries mature.

Leader and follower. Once a potential entrant joins an incumbent industry, the monopoly transitions into an oligopoly and both the competitive environment and innovation incentives change fundamentally. Within an oligopoly, the leader and followers are locked in a patent race in which each party's payoff depends on the other's actions and on the current state (m, n).

In an industry with a leader m steps ahead of n followers, the leader's normalized value v_{mn} satisfies

$$(\rho + g_N)v_{mn} = \max_{x \geq 0} \Big\{ \pi_{mn} - R(x) + x\big(v_{m+1,n} - v_{mn}\big) + nx_{-mn}\big(v_{m-1,n} - v_{mn}\big) + z\,\theta_{mn}\big(v_{mn+1} - v_{mn}\big) - \delta v_{mn} \Big\}. \quad (5)$$

In an oligopoly with gap m and n followers, the leader earns flow profits π_{mn} and chooses its innovation intensity x. A successful leader innovation arrives at the chosen rate and widens the gap to $m+1$, raising continuation value by $v_{m+1,n} - v_{mn}$. Followers innovate at common intensity x_{-mn}, so the leader faces a total catch-up intensity nx_{-mn}. A follower innovation narrows the gap to $m-1$ and reduces the leader's continuation value by $v_{mn} - v_{m-1,n}$. In addition, potential entrants arrive at rate z and join this incumbent industry with probability θ_{mn}, which increases the number of followers from n to $n+1$ and changes the leader's continuation value from v_{mn} to $v_{m,n+1}$. As before, exogenous exit arrives at rate δ.

Given follower behavior, the optimal innovation intensity of the leader is

$$x_{mn} = \left(\frac{\max\{v_{m+1,n} - v_{mn},\, 0\}}{\alpha} \right)^{\frac{1}{\psi - 1}}.$$

Equation (5) and the optimal policy highlight an escape-competition motive. The leader widens the technology gap to protect its rents. Two forces discipline this incentive. First, as the gap m widens, the leader's profit rises but the marginal value of widening the gap further falls, since catching up becomes increasingly difficult for followers and their competitive threat diminishes. Second, as the number of followers n increases, both the leader's profit and the marginal value of innovation decline, because each follower erodes the leader's market share and contributes to the total catch-up rate nx_{-mn}. Together these forces imply that the leader's innovation incentive is strongest in early-stage oligopolies with few followers and a small gap. It weakens as the industry either matures into a wide-gap duopoly or becomes crowded with followers.

Each follower's value v_{-mn} solves

$$(\rho + g_N)v_{-mn} = \max_{x \geq 0} \Big\{ \pi_{-mn} - R(x) + \big(x + (n-1)x_{-mn}\big)\big(v_{-m+1,n} - v_{-mn}\big) \\ + x_{mn}\big(v_{-m-1,n} - v_{-mn}\big) + z\,\theta_{mn}\big(v_{-mn+1} - v_{-mn}\big) - \delta v_{-mn} \Big\}. \quad (6)$$

A follower earns flow profits π_{-mn} and chooses an innovation intensity x. Because follower innovations are shared, the follower benefits not only from its own effort but also from other followers' effort. With total intensity $x + (n-1)x_{-mn}$, the follower group innovates and collectively moves one step closer to the leader, increasing the follower's continuation value from v_{-mn} to $v_{-(m-1),n}$. By contrast, leader innovation arrives at rate x_{mn} and moves the leader further ahead, reducing the follower's continuation value from v_{-mn} to $v_{-(m+1),n}$. Entry at rate $z\theta_{mn}$ adds an additional follower, tightening competition among followers, and exogenous exit arrives at rate δ.

The common innovation intensity of followers is

$$x_{-mn} = \left(\frac{\max\{v_{-m+1,n} - v_{-mn},\, 0\}}{\alpha} \right)^{\frac{1}{\psi - 1}}.$$

The follower's R&D problem highlights a tension between free-riding and catch-up motives. Because improvements diffuse among followers, each follower enjoys part of the benefit of peers' R&D, but the private return to investing in catch-up declines when the leader is already far ahead and when the follower tier is crowded.

When $m = 0$, firms are symmetric neck-and-neck competitors and the value of a representative firm v_{0n} satisfies

$$(\rho + g_N)v_{0n} = \max_{x \geq 0} \Big\{ \pi_{0n} - R(x) + x\big(v_{1n} - v_{0n}\big) \\ + n x_{0n}\big(v_{-1,n} - v_{0n}\big) + z\,\theta_{0n}\big(v_{0,n+1} - v_{0n}\big) - \delta v_{0n} \Big\}. \quad (7)$$

A firm's innovation at rate x moves it into a temporary leader position (state $m = 1$), while innovations by any of the n firms at total rate nx_{0n} can displace its current position. Entry adds an additional competitor at rate $z\theta_{0n}$, further intensifying business stealing.

This case isolates the classic mechanism that competition can raise innovation when firms are close to the frontier, because the prize from becoming the temporary leader is high. At the same time, the model disciplines this force through endogenous entry. As n rises, business stealing intensifies and the net private return to R&D can fall.

The structure of these Bellman equations is standard in step-by-step innovation models.

Leaders benefit from pulling ahead and suffer when followers catch up, while followers benefit from their own and their peers' innovation and lose when the leader moves further ahead. Entry and exit shift the number of followers in the industry and truncate the horizon, and the Poisson innovation structure ensures that all continuation values are linear in the normalized scale $\bar{Y}_t$.

Potential entrant. The entry decision is the central margin that links within-industry competition to the rate of industry creation. I assume that the cost of creating a new industry, κ, is drawn from a distribution G with support $[\underline{\kappa}, \bar{\kappa}]$, where $0 \leq \underline{\kappa} < \bar{\kappa} < v_1^M$ in equilibrium, so that even the highest-cost potential entrant would obtain strictly positive surplus from starting a new industry.

A potential entrant with cost draw κ compares two options. Joining an existing industry is cheap but competitive. The entrant adopts the lowest available technology and earns a follower value, either v_{01} if the industry is currently a monopoly or v_{-mn+1} if it is an oligopoly with gap m and n followers. Creating a new industry is costly but secures a temporary monopoly. The entrant pays $\kappa\bar{Y}$ and earns v_1^M, the value of a monopolist starting with a one-step lead over the public knowledge frontier. The entrant joins the incumbent industry if and only if the follower value exceeds the net value of creating a new industry, that is, if $v_{S+}^f \geq v_1^M - \kappa$, where v_{S+}^f denotes the relevant follower value in state S. The relative attractiveness of the two options depends critically on the state of the industry the entrant encounters. In young, competitive industries where the follower's catch-up probability is substantial, the follower value v_{S+}^f is high and direct entry is more attractive. In mature industries where technology gaps are wide and profits are compressed by crowding, follower values are depressed and more entrants find it worthwhile to bear the fixed cost of creating a new market instead.

These choices can be represented by indicator functions

$$\iota_\kappa^M = \begin{cases} 1 & \text{if } v_{01} \geq v_1^M - \kappa, \\ 0 & \text{otherwise,} \end{cases} \qquad \iota_{mn,\kappa} = \begin{cases} 1 & \text{if } v_{-mn+1} \geq v_1^M - \kappa, \\ 0 & \text{otherwise.} \end{cases}$$

Aggregating over κ using its distribution G yields the joining probabilities

$$\theta_M \equiv \int \iota_\kappa^M \, dG(\kappa), \qquad \theta_{mn} \equiv \int \iota_{mn,\kappa} \, dG(\kappa).$$

These are the probabilities that a potential entrant who meets a monopoly or an (m, n)-industry decides to join rather than create a new industry. They appear in the Bellman equations (4)-(7), closing the link between the entry decision and incumbent innovation. A higher joining probability in a given state means more frequent follower entry, which lowers leader values and dampens R&D investment by incumbents, while simultaneously reducing the flow of new-industry creation and slowing the economy's renewal rate.

The joining rules pin down state-contingent joining probabilities (θ_M, θ_{mn}). Whenever a potential entrant meets an incumbent industry in state ω, it creates a new industry with probability $1-\theta(\omega)$ and joins with probability $\theta(\omega)$. Since such meetings occur at rate z per incumbent industry, define the gross rate of variety creation per incumbent industry as

$$\gamma \equiv z\,\mathbb{E}[1-\theta(\omega)],$$

where the expectation is taken over the state of the industry encountered by the entrant, which coincides with the cross-sectional distribution in the stationary equilibrium characterized below. This object governs the evolution of the mass of industries with the exit rate δ and, together with ρ, determines the effective horizon for normalized firm values.

The next lemma summarizes key properties of the value functions and innovation rates in a stationary equilibrium.

Lemma 4.2 (Properties of value functions and innovation rates)**.** *In any stationary Markov perfect equilibrium, the normalized value functions and innovation rates satisfy:*

(i) For any normalized value v,

$$0 < v < \frac{1}{\rho+\gamma}.$$

(ii) The monopolist's value $v^M_{\hat{s}}$ increases in its technology lead $\hat{s}$, with

$$\lim_{\hat{s}\to\infty} v^M_{\hat{s}} = \frac{1+z\theta_M v_{01}}{\rho+\gamma+z\theta_M}.$$

The associated innovation rate $x^M_{\hat{s}}$ is strictly positive and decreases with $\hat{s}$, and satisfies $\lim_{\hat{s}\to\infty} x^M_{\hat{s}} = 0$.

(iii) For any $m \geq 0$, the values v_{mn}, v_{-mn} and their corresponding innovation rates x_{mn}, x_{-mn} converge to 0 as $n \to \infty$. For fixed n, we have

$$\lim_{m\to\infty} v_{mn} = \frac{1}{\rho+\gamma}, \qquad \lim_{m\to\infty} v_{-mn} = \lim_{m\to\infty} x_{mn} = \lim_{m\to\infty} x_{-mn} = 0.$$

(iv) There exists a finite upper bound $n^* \geq 1$ on the number of followers such that entry

into an industry with $n \geq n^*$ followers is never profitable. In particular, $\theta_{mn} = 0$ for all m and all $n \geq n^*$.

Lemma 4.2 characterizes how market structure shapes innovation incentives across different stages of the industry life cycle. Part (i) establishes that normalized firm values are finite and bounded above by $\frac{1}{\rho+\gamma}$, the present value of a perpetual unit profit stream discounted at the effective rate $\rho+\gamma$. The upper bound reflects two sources of discounting: the household's time preference ρ and the dilution of each industry's revenue share as new industries enter at rate γ. No firm can be worth more than this bound because flow profits are always below one.

Part (ii) describes how a monopolist's value and innovation effort evolve as it pulls further ahead of the fringe. A larger lead raises current profits, since $\pi^M_{\hat{s}}$ increases in $\hat{s}$, and therefore raises the monopolist's value. However, the marginal profit gain from one additional step diminishes. The incremental return to innovation, $v^M_{\hat{s}+1} - v^M_{\hat{s}}$, therefore shrinks with $\hat{s}$, and the monopolist's R&D effort declines even as its value rises. In the limit, as the monopolist's lead grows arbitrarily large, its value converges to $\frac{1+z\theta^M v_{01}}{\rho+\gamma+z\theta^M}$. At this limit, the marginal return to further innovation is zero while the entry threat still exists. This declining innovation effort of an established monopolist contributes to the growth slowdown within maturing industries.

Part (iii) reveals how competition shapes innovation incentives within oligopolies along two dimensions, the technology gap m and the number of followers n. As the number of followers grows holding m fixed, profits are competed away by crowding, follower values collapse to zero, and both leaders and followers cease to innovate. Too many competitors chasing the same rents destroy the private return to R&D for all parties. This is the extensive crowding effect. Along the other dimension, as the technology gap grows holding n fixed, the industry converges to a different limit. The leader's value approaches the upper bound $\frac{1}{\rho+\gamma}$, reflecting the near-monopoly profits of a dominant firm with an insurmountable lead, while follower values and innovation rates collapse to zero. Followers in such industries face an increasingly futile catch-up problem. The expected gain from a successful innovation, $v_{-m+1,n} - v_{-mn}$, shrinks as the gap widens, because closing one step still leaves the follower far behind. The leader, meanwhile, faces a negligible competitive threat and has little incentive to invest further. Wide-gap oligopolies thus converge to a low-innovation equilibrium that resembles a de facto monopoly, even though multiple firms remain active.

Part (iv) connects the crowding effect to the entry margin. Follower values fall to zero as n grows, so there exists a finite threshold n^* beyond which no potential entrant finds it worthwhile to join as an additional follower, even at the lowest possible market-creation cost $\underline{\kappa}$. Beyond n^*, all potential entrants who encounter that industry choose to create new industries instead, redirecting the flow of entrepreneurial activity toward the extensive margin, which keeps the relevant set of industry sizes finite in equilibrium.

4.3 Stationary Markov-Perfect Equilibrium

Given the optimal decisions described in the previous subsection, I restrict attention to stationary symmetric Markov-perfect equilibria and characterize their properties. Let $\mu_{\hat{s}}^M$ denote the share of monopoly industries whose leader is $\hat{s}$ steps ahead of the fringe, and let μ_{mn} denote the share of oligopoly industries with technology gap m and n followers. These shares satisfy

$$\sum_{\hat{s}=1}^{\infty} \mu_{\hat{s}}^M \;+\; \sum_{m=0}^{\infty} \sum_{n=1}^{n^*} \mu_{mn} \;=\; 1.$$

Let $\mu^M = \sum_{\hat{s}=1}^{\infty} \mu_{\hat{s}}^M$ denote the total share of monopolies. The gross rate of variety creation γ, defined in Section 4.2 as $z\,\mathbb{E}[1-\theta(\omega)]$, can now be expressed explicitly using the cross-sectional distribution:

$$\gamma \;=\; z\left[(1-\theta^M)\mu_M + \sum_{n=1}^{n^*} \sum_{m=0}^{\infty} (1-\theta^{mn})\mu_{mn}\right].$$

Here, z is the exogenous arrival rate of potential entrants, while γ is the endogenous equilibrium rate of new-industry creation. This expression formalizes the relationship between entry decisions and aggregate growth that was previewed in Section 4.2. The rate γ aggregates the fraction of potential entrants who choose to create new industries rather than join existing ones, weighted by the mass of industries in each state. A higher fraction of mature industries, where follower values are depressed and joining is unattractive, diverts more entrepreneurial activity toward industry creation and raises γ. The Kolmogorov forward equations that determine $\{\mu_{\hat{s}}^M, \mu_{mn}\}$ in a stationary equilibrium are given in the appendix.

To understand what shapes the stationary distribution, it is useful to trace the life cycle of a typical industry through the state space. Figure 1 summarizes the resulting flows. Every industry enters the economy as a fresh monopoly at $\hat{s} = 1$, contributing to the flow γ of newly created industries. From there, monopoly industries evolve along two margins. First, the monopolist climbs the quality ladder through successful innovation. At rate $x_{\hat{s}}^M$, the monopolist advances from $\hat{s}$ to $\hat{s}+1$, pulling further ahead of the fringe and raising its profit. Second, at rate $z\theta^M$, a potential entrant chooses to join the monopoly as a follower rather than create a new market. This converts the monopoly into a neck-and-neck duopoly at state $(0, 1)$ and initiates the oligopoly phase of the life cycle.

Once an industry enters the oligopoly phase, its evolution is governed by the interplay of leader and follower innovation and continued entry. The technology gap widens at rate

Figure 1: Flowchart of industry shares

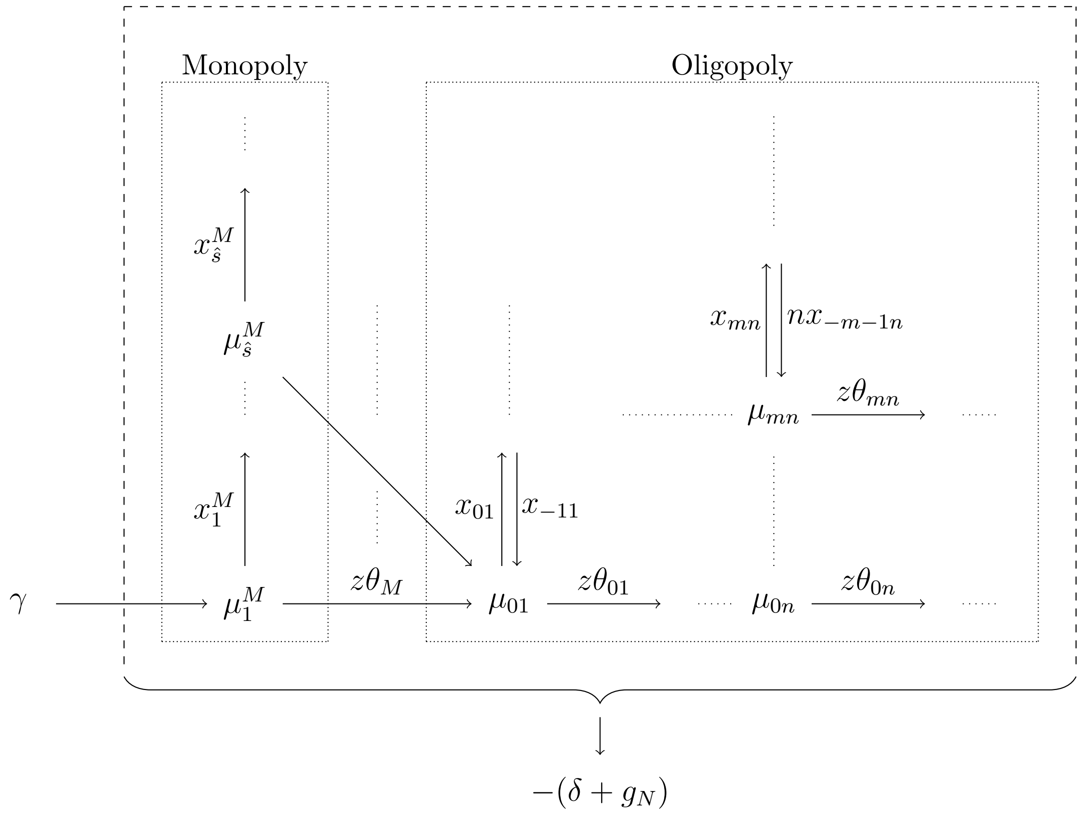

x_{mn} when the leader innovates and narrows at the total rate nx_{-mn} when any of the n followers successfully catches up. At the same time, new followers join at rate $z\theta_{mn}$, pushing the industry toward more crowded states. As Lemma 4.2 establishes, this crowding progressively depresses both leader and follower values and weakens innovation incentives for all parties. Joining probabilities eventually fall to zero once the industry reaches n^* followers. An industry therefore traverses a predictable arc. It is born competitive, passes through an oligopolistic phase of active innovation and entry, and gradually settles into a mature, low-innovation configuration before exiting at rate δ.

Each industry disappears at rate δ, and its relative share also declines at rate g_N due to the expanding mass of industries. The evolution of the total mass of industries N is governed by

$$g_N = \gamma - \delta,$$

the net rate of variety creation. In a stationary equilibrium, the distribution $\{\mu^M_{\hat{s}}, \mu_{mn}\}$ is time-invariant, meaning that the flows of industries into and out of each state exactly balance. The key aggregate state variables for growth are therefore the rate of variety creation γ and the distribution of gaps and followers summarized by $\{\mu^M_{\hat{s}}, \mu_{mn}\}$. I can now define the equilibrium formally.

Definition 4.3 (Equilibrium)**.** *A stationary and symmetric Markov-perfect equilibrium consists of values $\{v_{\hat{s}}^M, v_{mn}, v_{-mn}\}$, policies $\{x_{\hat{s}}^M, x_{mn}, x_{-mn}, \theta^M, \theta_{mn}\}$, an industry distribution $\{\mu_{\hat{s}}^M\}_{\hat{s}} \cup \{\mu_{mn}\}_{m,n}$, and scalars $\{g_N, \gamma\}$ such that all individual decisions solve the corresponding optimization problems and the laws of motion for industry shares are satisfied.*

With the equilibrium distribution in hand, I can characterize how the aggregate technology level evolves. Let $\bar{s}_t^M$ denote the average monopolist technology level and $\bar{s}_t^{cf}$ the average technology level of the associated competitive fringe firms. For oligopoly industries, let $\bar{s}_t^f$ denote the average lowest technology level and $\bar{s}_t^l$ the average highest technology level, and define the average technology gap among oligopolies as $\bar{m}$. The public knowledge stock can then be written as

$$s_t^K = \mu_M \bar{s}_t^M + (1-\mu_M)\bar{s}_t^f.$$

Lemma 4.4 (Properties of technology levels in a stationary equilibrium)**.** *The public knowledge stock grows at a constant rate*

$$\dot{s}^K = \gamma + \sum_{\hat{s}=1}^{\infty} \mu_{\hat{s}}^M x_{\hat{s}}^M + \sum_{n=1}^{n^*}\sum_{m=1}^{\infty} n x_{-mn}\mu_{mn}.$$

The average technology levels satisfy

$$\begin{aligned}
\gamma(\bar{s}_t^M - \bar{s}_t^{cf}) &= \gamma + \sum_{\hat{s}=1}^{\infty} \mu_{\hat{s}}^M x_{\hat{s}}^M, \\
\gamma(1-\mu_M)(\bar{s}_t^l - \bar{s}_t^f) &= \gamma(1-\mu_M)\bar{m} \\
&= \sum_{n=1}^{n^*}\left[(n+1)x_{0n}\mu_{0n} + \sum_{m=1}^{\infty} x_{mn}\mu_{mn}\right] - \sum_{n=1}^{n^*}\sum_{m=1}^{\infty} n x_{-mn}\mu_{mn}.
\end{aligned}$$

Lemma 4.4 characterizes the technology dynamics that underpin steady aggregate growth. The first result establishes that the public knowledge stock grows at a constant rate in a stationary equilibrium, combining three distinct sources of productivity advance. The first is variety creation at rate γ. Each new industry is born at a productivity level one step above the current public knowledge frontier, so the mere creation of new industries pushes up s^K even before any within-industry innovation occurs. The second and third terms capture within-industry innovation by monopolists and followers respectively. Monopolists push the knowledge frontier upward by innovating beyond the fringe, while followers contribute by catching up to the leader, which advances the lowest technology tier and thereby raises the public knowledge stock.

The second set of expressions in Lemma 4.4 links the average technology gaps to the balance of innovative activity across sectors. The average lead of monopolists over the fringe, $\bar{s}^M - \bar{s}^{cf}$, is maintained in steady state by opposing forces. New monopoly creation and monopolist innovation widen the average gap, while the expansion of the industry mass at rate γ continuously dilutes it. Similarly, the average oligopoly gap $\bar{m}$ is kept constant by the balance between leader innovation, which widens gaps, follower innovation, which narrows them, and the dilution by industry mass growth. These stationarity conditions are not mechanical accounting identities but reflect equilibrium restrictions. They require that the cross-sectional distribution of industries across states be consistent with the individual innovation incentives characterized in Section 4.2. These relationships underpin the aggregate growth decomposition in the next section.

5 Aggregate Growth and Industry Dynamics

This section derives the main theoretical results. I first decompose the aggregate growth rate into contributions from variety creation, frontier innovation, and catch-up innovation, each weighted by the stationary distribution characterized in Section 4. I then analyze the constrained efficiency of the variety creation margin and show that the decentralized economy creates too few new industries. Finally, I characterize how individual industries evolve over their life-cycles despite the stationarity of the cross-section.

5.1 Aggregate Growth Decomposition

Section 4 characterized the stationary cross-sectional distribution of industries, with constant shares in each life-cycle state and a well-defined rate of variety creation γ. I now show how the aggregate growth rate emerges from this stationary distribution. The central result is that steady aggregate growth arises even though individual industries follow non-stationary life-cycle paths. The stationary cross-section ensures that the average contribution of each growth margin remains constant over time.

Proposition 5.1 (Steady growth with non-stationary industry life-cycle)**.** *In any symmetric* *stationary Markov-perfect equilibrium, the constant growth rate can be expressed as:*

$$
\begin{aligned}
g &= \ln \lambda \left(\gamma + \sum_{\hat{s}=1}^{\infty} \mu_{\hat{s}}^{M} x_{\hat{s}}^{M} + \sum_{n=1}^{n^*} \sum_{m=1}^{\infty} n x_{-mn} \mu_{mn} \right) \\
&= \ln \lambda \left(\gamma + \sum_{\hat{s}=1}^{\infty} \mu_{\hat{s}}^{M} x_{\hat{s}}^{M} + \sum_{n=1}^{n^*} \left((n+1) x_{0n} \mu_0 + \sum_{m=1}^{\infty} x_{mn} \mu_{mn} \right) - \gamma (1 - \mu_M) \bar{m} \right) \\
&= \ln \lambda \cdot \dot{s}^K .
\end{aligned}
$$

The first line decomposes the growth rate into three distinct sources, each scaled by the step size $\ln \lambda$. The first term, γ, captures the contribution of variety creation. Each new industry is born with a productivity level one step above the current public knowledge frontier, so the flow of new industries directly raises the average technology level. With rate γ, a new monopolist with level $s^K + 1$ appears, but the s^K component is offset by the disappearance of incumbent industries at rate $\gamma = \delta + g_N$, leaving a net contribution of γ to growth. The second term aggregates the innovation effort of monopolists weighted by their share $\mu_{\hat{s}}^{M}$, capturing how incumbents push the technology frontier upward within existing monopoly industries. The third term captures catch-up innovation in oligopolies where followers innovate at the combined rate $n x_{-mn}$, advancing the lowest technology tier and thereby raising the public knowledge stock even without moving the frontier.

The second line rewrites the same expression in terms of the evolution of leaders' technology levels and the stationary gap $\bar{m}$, using Lemma 4.4. This alternative decomposition replaces the follower catch-up terms with leader innovation, net of the correction term $\gamma(1 - \mu_M)\bar{m}$. This correction arises because new industries replace exiting ones starting at $s^K + 1$. Only the s^K component of the exiting industry's technology is replaced by the first γ term. The above-lowest-level component of oligopoly leaders, which on average equals $\bar{m}$, simply disappears upon exit and is not replenished by the new entrant. The last equality simply notes that the sum of these contributions is, by definition, the rate of growth of the public knowledge stock $\dot{s}^K$.

Proposition 5.1 shows that aggregate growth is fully determined by industry dynamics. Variety creation, innovation in monopoly industries, and catch-up innovation in oligopolies all contribute to growth, and their relative importance depends on the stationary distribution of industries across life-cycle states. This structure implies that policies or parameter changes that shift the cross-sectional distribution affect aggregate growth through several margins simultaneously. Making entry into existing industries more or less attractive, for instance, changes the balance of all three growth components. I return to this point in Section 5.2.

The existence of a stationary equilibrium relies on two features of the model working in tandem. First, private innovation incentives are disciplined by diminishing marginal gains from pushing an already-large lead further and by congestion from an expanding follower base, which jointly make innovation and profits negligible in very mature, crowded industries. Second, endogenous industry creation continually replenishes the cross-section with new, innovation-intensive monopolies, offsetting the gradual maturation and exit of incumbent industries. Without the first feature, the cross-sectional distribution would unravel as gaps grow without bound. Without the second, the economy would eventually be populated only by mature, low-innovation industries and growth would stall.

5.2 Constrained efficiency of industry creation

Proposition 5.1 showed that aggregate growth depends on the stationary distribution of industries, which is shaped by entrants' choices between joining incumbent industries and creating new ones. A natural question is whether the equilibrium produces the right mix of industry creation and incumbent entry. To address this, I compare the equilibrium allocation with a constrained planner's benchmark.

The equilibrium entry decision in Section 4.2 is formulated as a binary choice for each potential entrant with a realized cost draw κ, between joining the encountered incumbent industry and creating a new one. To compare the equilibrium allocation with a social benchmark, it is useful to rewrite this decision in terms of state-specific cost thresholds. Let S index an industry state, where $S = M$ denotes a monopoly with technology gap $\hat{s}$ and $S = mn$ denotes an oligopoly with gap m and n followers. Let μ_S be the stationary share of industries in state S and let v^f_{S+} denote the follower value after joining that industry. For example, $v^f_{S+} = v_{01}$ if the industry is a monopoly and $v^f_{S+} = v_{-mn+1}$ if it is an oligopoly with gap m and n followers.

Before being randomly assigned to an incumbent industry and before observing its own cost draw κ, a potential entrant chooses a vector of thresholds $\{\hat{\kappa}^{CE}_S\}_S$ that solves

$$\max_{\{\hat{\kappa}_S\}_S} \sum_S \mu_S \left[\int_{\underline{\kappa}}^{\hat{\kappa}_S} \left(v_1^M - \kappa \right) dG(\kappa) + \int_{\hat{\kappa}_S}^{\overline{\kappa}} v^f_{S+} \, dG(\kappa) \right], \tag{8}$$

where G is the distribution of market initiation costs and v_1^M is the value of a new monopolist.[3] In each state S, the entrant draws a cost κ and compares two options. If $\kappa < \hat{\kappa}_S$, the entrant pays κ to create a new industry and earns $v_1^M - \kappa$. If $\kappa \geq \hat{\kappa}_S$, the entrant joins the incumbent industry and earns the follower value v^f_{S+}. The first-order condition for an

[3] I assume $\underline{\kappa} < v_1^M$ so that starting a new industry yields strictly positive surplus at low costs.

interior threshold in state S is

$$\mu_S G'(\hat{\kappa}_S^{CE}) \left[-\hat{\kappa}_S^{CE} + v_1^M - v_{S+}^f\right] = 0, \tag{9}$$

so that whenever $\mu_S G'(\hat{\kappa}_S^{CE}) > 0$ the equilibrium cutoff satisfies

$$\hat{\kappa}_S^{CE} = v_1^M - v_{S+}^f \approx \frac{\pi_1^M - \pi_{S+}^f}{\rho + \gamma}. \tag{10}$$

The approximation on the right replaces the full value difference with the flow profit differential $\pi_1^M - \pi_{S+}^f$ capitalized at the effective discount rate $\rho + \gamma$. The implied incumbent joining probability is

$$\theta_S \;=\; 1 - G\big(\hat{\kappa}_S^{CE}\big) \tag{11}$$

and it coincides with the probabilities defined in the entry problem above.

To assess the efficiency of this allocation, I consider a constrained planner who can intervene only in the variety creation margin. The planner does not directly control incumbent pricing or innovation decisions. Instead, the planner chooses a vector of thresholds $\{\hat{\kappa}_S^{SP}\}_S$ to maximize the representative household's lifetime utility, subject to the aggregate resource constraint and the induced law of motion for the industry distribution. Firms continue to optimize given the competitive environment they face. However, the planner recognizes that changing entry thresholds alters the stationary distribution of industries, which in turn affects firm values and innovation incentives through the equilibrium fixed-point.

For any given set of thresholds $\{\hat{\kappa}_S\}_S$, the induced entry probabilities $\{\theta_S\}_S$ determine a unique stationary distribution over industry states and a corresponding balanced growth path for consumption. Let $\mathcal{R}(\{\hat{\kappa}_S\}_S)$ denote the ratio of incumbent R&D expenditure to aggregate output and $\mathcal{K}(\{\hat{\kappa}_S\}_S)$ the ratio of market initiation costs to output as defined in Appendix B.3, and $g(\{\hat{\kappa}_S\}_S)$ the implied aggregate growth rate. With log preferences, the planner's welfare can be written as

$$W\big(\{\hat{\kappa}_S\}_S\big) = \frac{1}{\rho} \ln\big(1 - \mathcal{R}\big(\{\hat{\kappa}_S\}_S\big) - \mathcal{K}\big(\{\hat{\kappa}_S\}_S\big)\big) + \frac{1}{\rho^2}\, g\big(\{\hat{\kappa}_S\}_S\big), \tag{12}$$

where the first term captures the level of net consumption on the balanced growth path and the second term captures the contribution of long-run growth to lifetime welfare. The coefficient $1/\rho^2$ arises from log preferences and reflects that a permanent increase in the growth rate affects consumption at all future dates. Because these gains accumulate linearly over time and are discounted at rate ρ, a one-unit increase in g raises welfare by $1/\rho^2$.

The constrained planner's first-order condition for the cutoff in state S can be written as

$$\frac{z}{\rho}\,\mu_S G'\big(\hat{\kappa}_S^{SP}\big)\left[-\frac{1}{1-\mathcal{R}-\mathcal{K}}\,\hat{\kappa}_S^{SP}+\frac{1}{\rho}\ln\lambda\right]+\Phi_S=0, \tag{13}$$

where all objects are evaluated at $\{\hat{\kappa}_S^{SP}\}_S$. The term $\frac{z}{\rho}\mu_S$ is the mass of potential entrants throughout history who encounter an industry in state S, and $G'(\hat{\kappa}_S^{SP})$ weights the marginal entrant at the cutoff. Inside the brackets, $-\frac{1}{1-\mathcal{R}-\mathcal{K}}\hat{\kappa}_S^{SP}$ is the utility cost of the market initiation expenditure, evaluated at the shadow value of consumption. The term $\frac{1}{\rho}\ln\lambda$ is the welfare gain from permanently raising the knowledge level by one quality step. The term Φ_S collects the marginal welfare effect of changing industry composition and incumbent innovation rates when the vector of thresholds is perturbed.

Comparing the planner's first-order condition (13) with the private entrant's optimality condition (9) reveals the sources of inefficiency in the decentralized equilibrium. The private entrant equates the market initiation cost $\hat{\kappa}_S^{CE}$ to the value differential $v_1^M - v_{S+}^f$, which is approximately $\frac{\pi_1^M-\pi_{S+}^f}{\rho+\gamma}$. The planner instead weighs the shadow cost $\frac{1}{1-R-K}\hat{\kappa}_S^{SP}$ against the social benefit $\frac{\ln\lambda}{\rho}$, plus the composition adjustment Φ_S. Four distinct forces drive a wedge between these two cutoffs. Three operate through the direct terms inside the brackets of (13), while a fourth arises from the general equilibrium response of the industry distribution captured by Φ_S. I discuss each in turn.

First, there is a *resource-cost externality* distorting the cost valuation. Private firms treat the market initiation cost as an expense in output units, whereas the planner applies the shadow value $\frac{1}{1-\mathcal{R}-\mathcal{K}}$ to account for the aggregate resource constraint. Higher aggregate investment in R&D and market creation raises the marginal utility of consumption, making the utility cost of funding variety creation strictly higher than its output-unit cost.

Second, there is a *knowledge externality*. Private entrants value variety creation through the flow of appropriable rents, $\pi_1^M - \pi_{S+}^f$, whereas the planner values the permanent expansion of the knowledge stock by one quality step λ. Since $\ln\lambda > \pi_1^M = 1-\lambda^{-1}$ for $\lambda > 1$, the private rent necessarily falls short of the social return to variety creation. Potential entrants undervalue the contribution of new industries to the aggregate knowledge stock because they cannot appropriate the full productivity gain.

Third, there is a *creative-destruction externality* that amplifies this wedge through the discount rate. The private discount rate, $\rho+\gamma$, is strictly larger than the social discount rate, ρ, because new industries arrive at rate γ and erode incumbent market shares. Private firms therefore discount future profits more heavily than is socially warranted. The planner discounts knowledge accumulation only at the time preference rate ρ, recognizing that the productivity gain $\ln\lambda$ persists beyond the life of the firm.

Finally, there is an *equilibrium adjustment effect*. The term Φ_S embeds the general equilibrium response of values, policies, and industry shares when the entire vector of thresholds is perturbed. A local change in $\hat{\kappa}_S$ affects not only the flow of entrants into state S but also the induced evolution of other states through the stationary distribution and the aggregate resource constraint. This adjustment further shifts the planner's preferred cutoff away from the private one. In the next section, I evaluate these conditions in the calibrated economy and quantify the resulting entry wedge.

5.3 Industry Life-Cycle Dynamics

The entry decision that determines aggregate growth and shapes the efficiency of industry creation also governs how individual industries evolve over their life-cycles. The stationary cross-sectional distribution characterized in Section 4.3 is not sustained because industries individually reach a steady state. Rather, the maturation of incumbent industries is continuously offset by the creation of new ones and the exit of old ones. The following results characterize the implied dynamics of surviving industries and connect them to the empirical patterns documented in Section 2.

Proposition 5.1 characterized steady aggregate growth using stationary cross-sectional objects. A stationary cross-sectional distribution does not imply that old surviving industries are representative of the cross-section. Because industries are continually replenished by newborn oligopolies, surviving cohorts are selectively older and more mature. Cohort-level technology gaps can therefore drift upward over the life-cycle even while the cross-sectional mean gap remains constant.

Let $(\tilde{m}, \tilde{n})$ denote an oligopoly state drawn from the stationary cross-sectional distribution μ conditional on being an oligopoly, and let $\tilde{m}_\Delta$ denote the technology gap after a short interval Δ under the within-oligopoly transition dynamics, starting from the cross-sectional draw $(\tilde{m}, \tilde{n})$. Let $\hat{m}_a$ denote the realized technology gap of an oligopoly industry at age a, where age is measured from the first time the industry has at least one follower, so that $a = 0$ corresponds to a newborn oligopoly. Define

$$f(a) \equiv \mathbb{E}[\hat{m}_a \mid \text{alive at age } a] \, .$$

Let $U \sim \text{Exp}(\gamma)$ denote the age of a randomly sampled surviving oligopoly from the stationary cross-section. Then

$$\bar{m} = \mathbb{E}[f(U)], \qquad \bar{m}_\varepsilon = \mathbb{E}[f(U) \mid U > \varepsilon].$$

That is, $\bar{m}_\varepsilon$ is the average technology gap among surviving oligopolies older than ε.

Assumption 5.1. *For sufficiently mature surviving oligopolies, leader innovation weakly dominates follower catch-up on average.*

Proposition 5.2 (Cross-sectional and cohort technology-gap dynamics)**.** *In a stationary Markov-perfect equilibrium with stationary cross-sectional industry distribution μ, oligopoly technology gaps satisfy:*

(i) In the stationary cross-section, technology gaps tend to widen on average at the industry level:

$$\lim_{\Delta \downarrow 0} \frac{1}{\Delta} \mathbb{E}\big[\tilde{m}_\Delta - \tilde{m}\big] = \gamma \bar{m} > 0.$$

(ii) Older surviving oligopolies have above-average gaps: for sufficiently small $\varepsilon > 0$,

$$\bar{m}_\varepsilon > \bar{m}.$$

(iii) Under Assumption 5.1, $f(a)$ is weakly increasing for sufficiently large ages.

Proposition 5.2 makes clear that stationarity in this model is an aggregate and cross-sectional property, not an industry-level one. Part (i) shows that even when an oligopoly is drawn from the stationary cross-section, its technology gap tends to widen on average under the within-industry dynamics. Thus, the stationary distribution is not sustained because industries themselves become individually stationary. Rather, it is sustained because exit and the creation of new industries continuously replenish the cross-section with young oligopolies.

Part (ii) is the cohort counterpart of this observation. The stationary cross-section places disproportionate weight on young oligopolies, which begin with low technology gaps. As a result, once a sufficiently small neighborhood of newborn surviving oligopolies is excluded, the average technology gap exceeds the stationary cross-sectional mean. A stationary cross-section is therefore not representative of surviving cohorts. This result does not yet imply a general monotone age profile, however, which requires the additional condition in Part (iii).

Part (iii) strengthens the result to a monotone age profile under the additional condition of Assumption 5.1 on the balance of leader and follower innovation in mature industries. This stronger cohort result again relies on endogenous industry creation. New industries are continually born with low technology gaps, which replenishes the cross-section with young oligopolies even as surviving cohorts become progressively more mature. In the calibrated economy, I verify numerically that Assumption 5.1 is satisfied and that $f(a)$ is strictly increasing over the relevant range of ages. As technology gaps widen along surviving cohorts, the marginal return to innovation falls for both leaders and followers, entry becomes less attractive, and leadership turnover becomes rare.

Since follower counts do not decrease in the baseline model, Proposition 5.2 implies that older surviving oligopolies tend to have both larger technology gaps and weakly larger follower bases. Larger gaps and more followers lower relative prices and raise industry quantities, so the gap dynamics developed here connect cohort maturation directly to the price and quantity implications of Proposition 5.3. The next proposition collects these life-cycle implications and connects them to the empirical regularities documented in Section 2.

To state the life-cycle implications, define the following industry-level variables. Let $P_{mn}(s^f)$ denote the industry price of an oligopoly with technology gap m, n followers, and follower technology level s^f. Let $Y_{mn}(s^f)$ denote the corresponding industry output.

Proposition 5.3 (Properties of industry life-cycle). *In a stationary Markov perfect equilibrium, conditional on industry survival:*

(i) Monopoly quality growth slows:

$$\lim_{\Delta \downarrow 0} \frac{1}{\Delta} \mathbb{E}\big[\ln q^M_{a+\Delta} - \ln q^M_a \,\big|\, \text{still monopoly, alive}\big] \searrow 0 \qquad \text{as } a \to \infty.$$

(ii) Mature oligopoly states have lower relative prices and higher quantities: For each $n \in \{1, \dots, n^*\}$, there exists $M_P(n) < \infty$ such that

$$\frac{P_{m+1,n}(s^f)}{w} \leq \frac{P_{mn}(s^f)}{w}, \qquad Y_{m+1,n}(s^f) \geq Y_{mn}(s^f) \qquad \text{for all } m \geq M_P(n).$$

Moreover, for each m and each $n \in \{1, \dots, n^* - 1\}$,

$$\frac{P_{m,n+1}(s^f)}{w} \leq \frac{P_{mn}(s^f)}{w}, \qquad Y_{m,n+1}(s^f) \geq Y_{mn}(s^f).$$

(iii) Under Assumption 5.1, for sufficiently large ages, surviving oligopoly cohorts move toward lower-price, higher-quantity, lower-innovation, lower-entry, and lower-turnover states.

Proposition 5.3 summarizes the industry-level forces that generate the model's life-cycle patterns. Steady aggregate growth does not arise because individual industries are stationary. Instead, industries move through non-stationary stages, while the aggregate economy remains on a balanced growth path because new industries are continually created and the cross-sectional distribution of industries is stationary.

Part (i) characterizes the monopoly stage. As a surviving monopolist extends its lead over the fringe, the marginal gain from further innovation declines. Monopoly quality growth therefore slows over time until the industry eventually attracts entry and transitions into oligopoly.

Part (ii) characterizes mature oligopoly states. For a given follower base, a sufficiently large technology gap lowers the industry price relative to the wage and raises industry quantity. For a given technology gap, a larger follower base has the same effect through stronger within-industry competition. These state-level properties provide a micro-foundation for one of the stylized facts emphasized in Section 2, namely that industry prices tend to fall while industry quantities rise over the life-cycle, even though the quantities of individual products may decline as they are replaced by newer varieties (Gort and Klepper, 1982; Filson, 2001; Argente, Lee and Moreira, 2024).

Part (iii) translates the cohort gap dynamics of Proposition 5.2 into broader life-cycle implications. As surviving oligopoly cohorts age, they move toward lower-price, higher-quantity, lower-innovation, lower-entry, and lower-turnover states. The industry life-cycle evidence confirms this pattern. Early phases are marked by rapid innovation and entry, while later phases are characterized by consolidation and lower technological activity (Abernathy and Utterback, 1978; Gort and Klepper, 1982; Klepper, 1996; Filson, 2001; Jovanovic and Wang, 2025).

The economy therefore does not converge to a representative mature industry. Instead, the aggregate growth is sustained by the continual creation, evolution, and exit of industries. It is precisely the entry of new, innovation-intensive industries that prevents the cross-section from being dominated by mature, low-growth sectors. This positive mechanism, together with the constrained efficiency analysis in Section 5.2, provides the foundation for the quantitative analysis in Section 6.

6 Quantitative Analysis

This section disciplines the theoretical framework quantitatively. I calibrate a stationary equilibrium to U.S. data from 1982–2011, and then use the calibrated economy to address three questions. First, does the model generate non-stationary industry life-cycles that are qualitatively consistent with the empirical patterns documented in Section 2? Second, does the decentralized economy create too few new industries, as the constrained efficiency analysis in Section 5.2 predicts? Third, how do the composition and intensity of entrepreneurial entry affect aggregate growth? And what does the cross-sectional structure of industries imply for the design of innovation policy?

Table 1: Calibrated Parameters

Externally calibrated		
ρ	0.0228	Rate of time preference (10-year Treasury rate − TFP growth)
δ	0.1500	Exit rate (brand exit rate (Broda and Weinstein, 2010))
Internally calibrated		
σ	3.8961	Elasticity of substitution
λ	1.0351	Innovation step size
α	0.0897	Multiplier of R&D cost function
ψ	1.2917	Curvature of R&D cost function
$\hat{k}$	0.0872	Mean of market initiation cost divided by baseline v_1^M
z	0.3115	Potential entrant arrival rate

For the quantitative analysis, I parameterize the distribution of market initiation costs as follows. Let κ denote the cost of creating a new industry, and write

$$\kappa = k\,\tilde{v}_1^M,$$

where k is drawn from a transformed beta distribution with support $[0, 0.9]$ and $\tilde{v}_1^M$ is the value of a monopolist in the baseline equilibrium.[4] In comparative statics, I vary the mean of this distribution while holding its standard deviation fixed at 0.1.

6.1 Calibration

I calibrate eight parameters to reflect the U.S. economy over the 30-year period starting in 1982. Table 1 reports the parameter values. Two parameters are chosen externally. The rate of time preference ρ is set to 0.0228, the difference between the average 10-year Treasury rate and average TFP growth, where TFP is measured using the utilization-adjusted series from Fernald (2014). The industry exit rate δ is set to 0.15, based on the median 1-year brand exit rate in Broda and Weinstein (2010).

The remaining six parameters, $\{\sigma, \lambda, \alpha, \psi, \hat{k}, z\}$, are estimated jointly (details are provided in the appendix), targeting the empirical moments in Table 2. While all parameters simultaneously affect all moments in general equilibrium, the following discussion identifies which data moments primarily pin down each parameter.

The elasticity of substitution σ governs the intensity of static price competition, while the market initiation cost parameter $\hat{k}$ governs the extensive margin of entry. Together they determine the level and dispersion of within-industry competition. I identify σ and $\hat{k}$ by

[4] This normalization ties the support of market initiation costs to the value of a new monopolist, which is convenient for comparing cutoffs across equilibria.

matching the average and standard deviation of cost-weighted markups. I compute these markup moments using Compustat data and the production approach from De Loecker, Eeckhout and Unger (2020) with input weights, which is more comparable to the product-level environment of the model.

Table 2: Targeted moments

	Data	Model	Source
Average TFP growth	0.0091	0.0091	Fernald (2014)
Average cost-weighted markup	1.3667	1.3717	Compustat, De Loecker, Eeckhout and Unger (2020)
S.D. of cost-weighted markup	0.2904	0.2872	Compustat, De Loecker, Eeckhout and Unger (2020)
Average R&D/GDP	0.0255	0.0252	NCSES (NSF), 2025
Average innovation frequency	0.5195	0.5198	Cai and Tian (2024) (1981–2014)
1-year median brand entry rate	0.1618	0.1586	Broda and Weinstein (2010) (1994, 1999–2003)

The R&D cost parameters (α, ψ) govern the cost of innovation for incumbents. I identify them by targeting the average R&D-to-GDP ratio, from NCSES (National Center for Science and Engineering Statistics , NCSES), and the average innovation frequency, from Cai and Tian (2024) covering 1981–2014. In the model, both incumbent R&D spending and market creation costs are counted as R&D, since both represent forms of innovation.

The innovation step size λ determines the return to innovation and is the primary driver of the aggregate growth rate. The parameter is calibrated to match the average TFP growth rate. Finally, z governs the frequency of potential entrant shocks and is identified by the median brand entry rate.

6.2 Industry Life-Cycle Patterns

The model's central mechanism is that endogenous industry creation generates non-stationary life-cycles whose cross-sectional composition determines aggregate growth. In the calibrated economy, industries follow precisely this pattern. They start small and highly innovative, then grow in size while innovation activity gradually slows. The profiles in Figure 2 show that the calibrated economy produces life-cycle dynamics that are qualitatively consistent with the empirical evidence discussed in Section 2. Each panel shows the expected change in an industry variable over a 50-year period, conditional on survival. Details of the calculation are provided in the appendix.

The top-left panel plots changes in output relative to the initial level. I report two measures. *Unadjusted output* (dash-dotted line) corresponds to the industry-level aggregator in the model and is best interpreted as the output of a narrowly defined market, close to a group of brands that compete directly. *Variety-adjusted output* (solid line) adjusts

Figure 2: Industry Life-Cycle Patterns

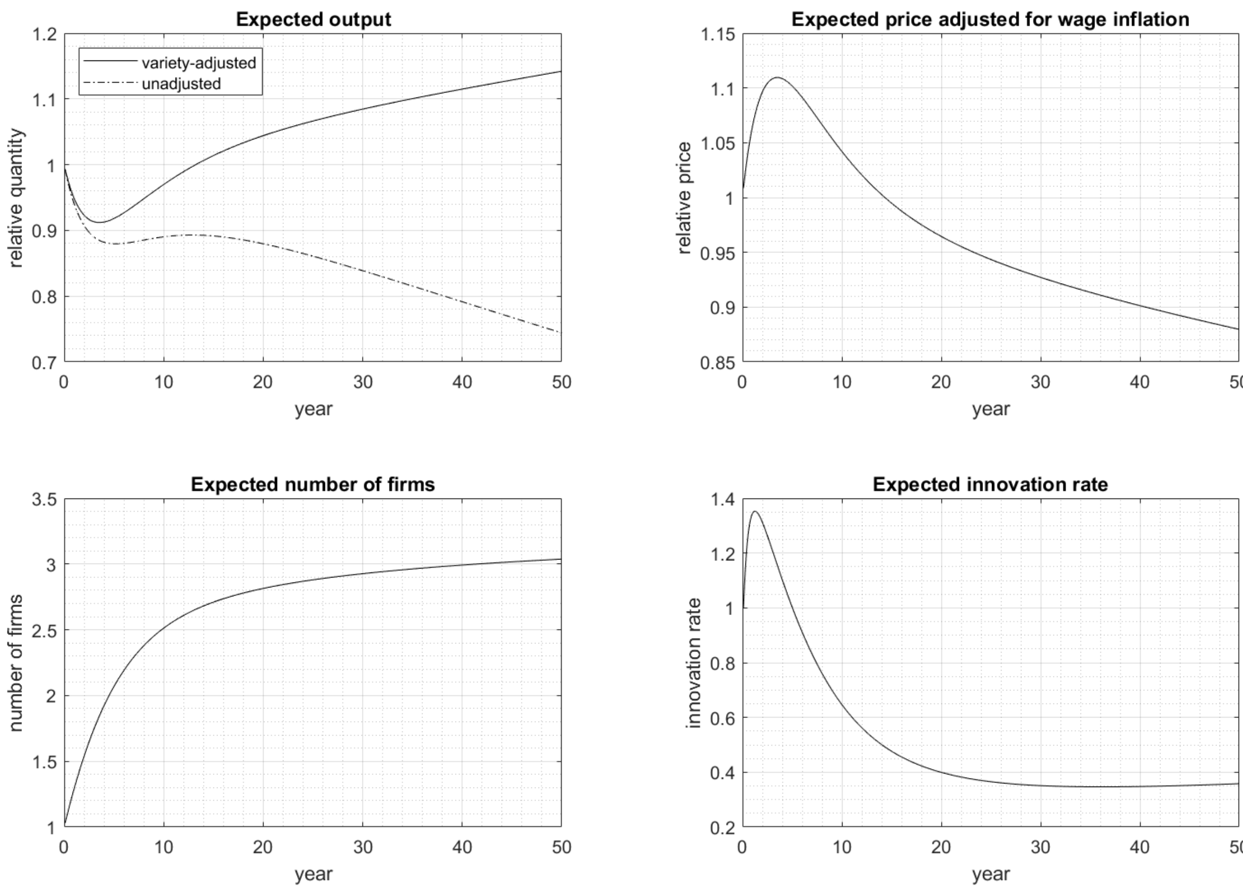

for changes in effective market size due to between-industry business-stealing effects from variety creation. In practice, new varieties often do not immediately generate new industry classifications, so business stealing is more salient at the brand level than at the industry level. The unadjusted series therefore aligns more closely with brand-level dynamics, while the adjusted series is more comparable to industry-level outcomes. As documented by Gort and Klepper (1982), Jovanovic and MacDonald (1994), and Filson (2001), industry-level output tends to rise over time, whereas brand-level output tends to decline as newer products capture market share. Argente, Lee and Moreira (2024) emphasize this latter pattern. The model reproduces both trends. Variety-adjusted output rises over the life-cycle, while unadjusted output eventually declines as new varieties compete away demand.

The top-right panel shows the expected change in industry prices, adjusted for wage inflation, relative to the initial level. Using the industry price index defined in Section 3 and dividing by the wage to account for aggregate inflation, industry prices fall over the life-cycle as innovation reduces unit costs and competition intensifies, consistent with the evidence.

The bottom panels plot the expected number of firms and the expected innovation rate over time. Firm counts increase while innovation rates decline, consistent with the stylized patterns documented by Abernathy and Utterback (1978), Filson (2001), and Jovanovic and Wang (2025). This trajectory captures the early expansion with frequent entry, followed by consolidation and lower innovation in more mature industries. The model does not feature an explicit shakeout phase with declining firm counts as in Klepper (1996), but such a pattern could be generated by introducing a fixed operating cost. As technology gaps widen and follower profits fall, some firms would optimally exit.

These age profiles are not directly targeted in the calibration in Table 2, so their qualitative consistency with the empirical patterns provides independent support for the model's central mechanism. Endogenous industry creation generates non-stationary life-cycles whose cross-sectional composition determines aggregate growth. In addition, the calibrated economy satisfies Assumption 5.1: for sufficiently mature surviving oligopolies, leader innovation weakly dominates follower catch-up on average, and the cohort expected gap function $f(a)$ is strictly increasing over the relevant range of ages. This confirms that the monotone life-cycle properties established in Proposition 5.3 (iii) hold in the quantitative model, not only as theoretical possibilities.

6.3 Constrained Efficiency of Calibrated Economy

I now use the constrained planner framework of Section 5.2 to assess the efficiency of variety creation in the calibrated economy. The goal is to quantify the wedge between the decentralized entry cutoffs $\{\hat{\kappa}_S^{CE}\}_S$ and the planner's cutoffs $\{\hat{\kappa}_S^{SP}\}_S$ that maximize welfare subject to the equilibrium laws of motion.

The planner's first-order condition for the cutoff in state S, equation (13), combines three components. The first is a congestion effect through the shadow value of aggregate consumption $1/(1-\mathcal{R}-\mathcal{K})$. The second is a knowledge effect through the term $(\ln \lambda)/\rho$. The third is a composition term Φ_S that captures how reallocating entrants across states affects the cross-sectional distribution of industries and incumbent innovation incentives. In this subsection, I evaluate these objects at the competitive equilibrium.

The congestion effect is modest but non-trivial. By construction, the planner attaches a higher utility cost to market initiation than private entrants. The shadow value of aggregate consumption is scaled by $1/(1-\mathcal{R}-\mathcal{K}) > 1$, so any given resource cost is magnified relative to the private valuation. In the calibrated economy, the planner's effective valuation of the market initiation cost is 2.59% higher than the private valuation for a given cutoff $\hat{\kappa}$.

The knowledge component pushes in the opposite direction and dominates. Although the private value v_1^M exceeds the simple bound $\pi_1^M/(\rho+\gamma)$ because it includes the option value of future innovation, the net private incentive $v_1^M - v_{S+}^f$ remains strictly below the social benefit $(\ln\lambda)/\rho$ for any state with positive mass ($\mu_S > 0$). Ignoring higher-order composition effects, the planner's cutoff would satisfy

$$\hat{\kappa}^{SP} = \frac{1-\mathcal{R}-\mathcal{K}}{\rho}\ln\lambda.$$

Under the current calibration, this expression evaluates to 1.47, which is strictly above the maximum cost $\overline{\kappa} = 0.57$ in the support of G. Conditional on holding innovation rates and industry shares fixed, it is therefore socially desirable to bear higher market initiation costs and create more varieties than in the competitive equilibrium.

To quantify the full equilibrium wedge, I compute the derivative of welfare with respect to the state-contingent thresholds at the competitive equilibrium. Evaluating the planner's first-order condition at the decentralized cutoffs and numerically approximating the derivatives, I find that they are positive for all economically relevant states ($\mu_S > 2\%$). For any state with substantial mass, variety creation is chosen less frequently than the constrained optimum. When I recompute these derivatives while keeping innovation rates and industry shares fixed at their competitive values, so that only the direct externalities enter and the composition terms Φ_S are suppressed, they remain positive for all active states, consistent with the comparison above between $v_1^M - v_{S+}^f$ and $(\ln\lambda)/\rho$. The equilibrium adjustment effect therefore does not overturn the sign of the wedge.

The calibrated economy thus exhibits too little variety creation relative to the constrained planner's benchmark. Potential entrants are too willing to enter existing oligopolies to chase short-term business-stealing rents, and too reluctant to pay the fixed costs to open new markets. The policy experiments in Section 6.5 explore the quantitative importance of this wedge by comparing two instruments that operate on different margins of the innovation process, subsidies aimed at variety creation and subsidies aimed at incumbent R&D.

6.4 Comparative Statics

To understand how the cross-sectional distribution of industries transmits parameter changes into aggregate growth, I conduct comparative statics with respect to parameters governing industry creation and competition. I first vary the potential entrant arrival rate z and the market initiation cost parameter $\hat{k}$, which control the intensity and cost of entry. I then vary the elasticity of substitution σ, which governs the strength of within-industry competition.

The left panel of Figure 3 shows how aggregate growth and its components respond to

Figure 3: Comparative Statics: Changing z and $\hat{k}$

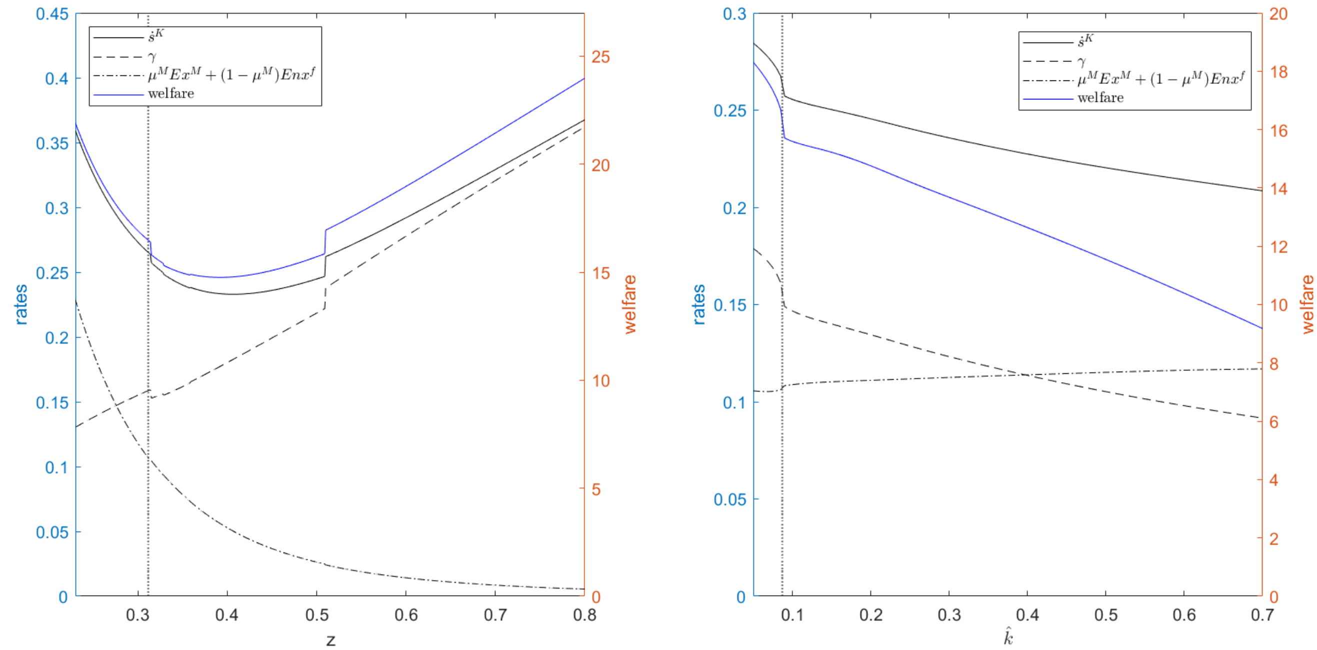

changes in z. The net result is a U-shaped relationship between z and aggregate growth. As z increases, the variety creation rate γ rises, but the average innovation rates of monopolists and oligopoly followers decline. More potential entrants generate more new industries, but they also reduce incumbent profits through both between- and within-industry business-stealing effects, lowering the expected value of innovation and crowding out R&D investment by incumbent firms. When z is low, increasing it reduces growth because the crowding-out effect dominates. Further increases beyond a threshold raise growth as the extensive-margin benefit from greater variety creation outweighs the loss from reduced incumbent innovation. The discrete jumps visible in these series reflect threshold switches in state-contingent entry conditions induced by the bounded support of market initiation costs, and are discussed in Appendix **??**

The right panel varies the mean market initiation cost $\hat{k}$. Reducing $\hat{k}$ uniformly lowers the cost of starting new industries and monotonically increases the growth rate. As $\hat{k}$ falls, the variety creation rate γ rises steadily. Incumbents' R&D incentives decline, but the magnitude is smaller than under changes in z, because the within-industry business-stealing channel is weaker. Leaders are affected more than followers, as the value of market leadership falls when the economy is filled with new and differentiated product lines. However, the overall reduction in R&D is muted by a shrinking average technology gap and a smaller

Figure 4: Comparative Statics: Changing σ

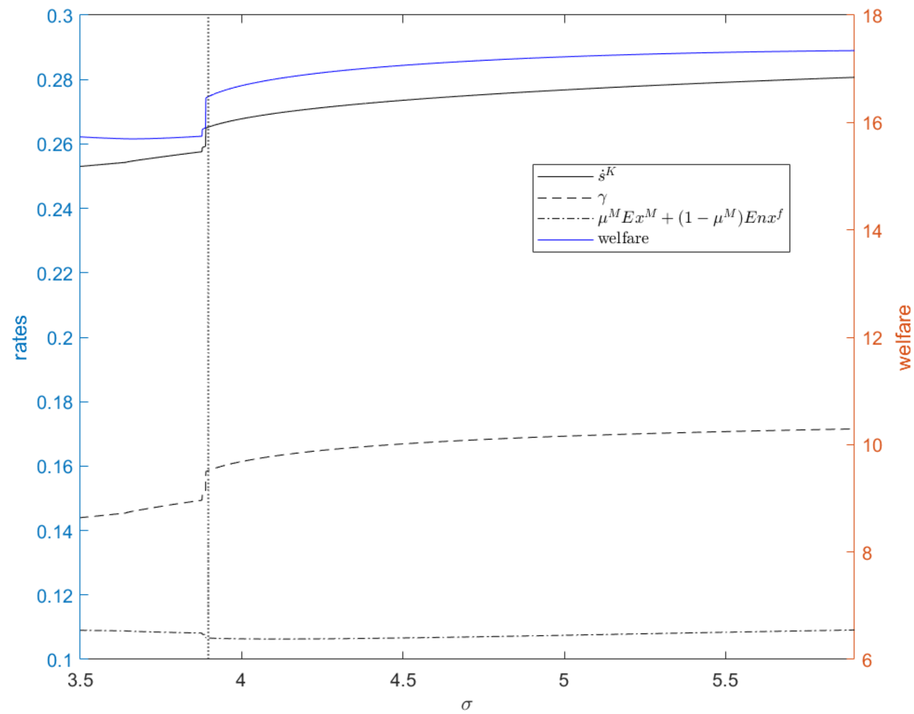

oligopoly share. Growth rises despite lower incumbent innovation because the extensive margin dominates: more frequent variety creation and the entry of younger, more competitive industries more than compensate for the decline in within-industry R&D.

These experiments also have a direct policy interpretation. The U-shaped relationship between z and growth cautions against the view that more entrepreneurs are always the right response to low growth. In the model, indiscriminately increasing the mass of potential entrants can move the economy into a region where additional entry mainly intensifies business stealing in existing industries and crowds out high-value incumbent R&D, with only limited gains in variety creation. What matters for long-run growth is not just the quantity of startups but their composition. Growth is highest when entry takes the form of new-industry creation rather than incremental entry into incumbent markets. Reducing the effective initiation cost $\hat{k}$ is a more targeted way to tilt entry toward new industries and raise growth, but it is also harder to implement in practice. The parameter $\hat{k}$ reflects deep features of the business environment, including regulation, fixed setup costs, and institutional barriers, that are difficult to change quickly.

Figure 4 reports comparative statics for the elasticity of substitution σ. The effect on incumbent innovation is modest and sign-ambiguous, reflecting the importance of industry composition. In industries with small technology gaps, stronger competition encourages both leaders and followers to invest more aggressively. In industries with large gaps, higher σ discourages follower innovation because catching up becomes less attractive. In the calibrated economy, the average incumbent innovation rate decreases slightly when σ is low and increases slightly when σ is high, but the magnitude of these changes is small.

The effect on variety creation is more pronounced. As it becomes less attractive to enter existing industries as a follower, potential entrants are more likely to initiate new industries instead. The variety creation rate γ rises with σ, and this extensive-margin effect dominates, so the aggregate growth rate increases. This pattern reveals a complementarity between pro-competitive product-market policies and technological diversification. Tougher within-industry competition fosters growth not only by stimulating "escape competition" innovation by incumbents, but also by accelerating structural change through variety creation.

The muted and sign-ambiguous response of incumbent innovation to σ echoes the findings of Aghion et al. (2005), who show that the competition-innovation relationship depends on firms' distance to the technological frontier. In this environment, however, the main action comes from the extensive margin: stronger competition within industries pushes potential entrants away from mature sectors and toward the creation of new industries, raising γ and, through this channel, the aggregate growth rate. The analysis thus adds a complementary mechanism to this literature, one that operates through the cross-sectional distribution of industries over life-cycle states.

6.5 Policy Experiments: Variety Creation versus Incumbent R&D

The constrained efficiency analysis in Section 6.3 established that the calibrated economy produces too little variety creation. I use that wedge here to compare two policy instruments that operate on different margins of the innovation process. Both policies are financed through lump-sum taxes on the representative household:

1. *Policy 1 (P1): Rewarding successful R&D investment.* A firm receives a subsidy equal to $b_r \Delta V$ upon a successful innovation, where $b_r \in [0, 1)$ is the subsidy rate and ΔV is the value increase resulting from the innovation.

2. *Policy 2 (P2): Rewarding variety creation.* A potential entrant receives a subsidy equal to $b_i V_1^M$ if it chooses to start a new industry, where $b_i \in [0, 1)$ is the subsidy rate and V_1^M is the value of a new monopolist.

Policy 1 directly targets incumbent R&D, aiming to boost technological improvement within existing industries. Policy 2 instead targets the extensive margin by increasing the returns to successful variety creation.

Figure 5: Policy Comparison: Rewarding Incumbent R&D and Variety Creation

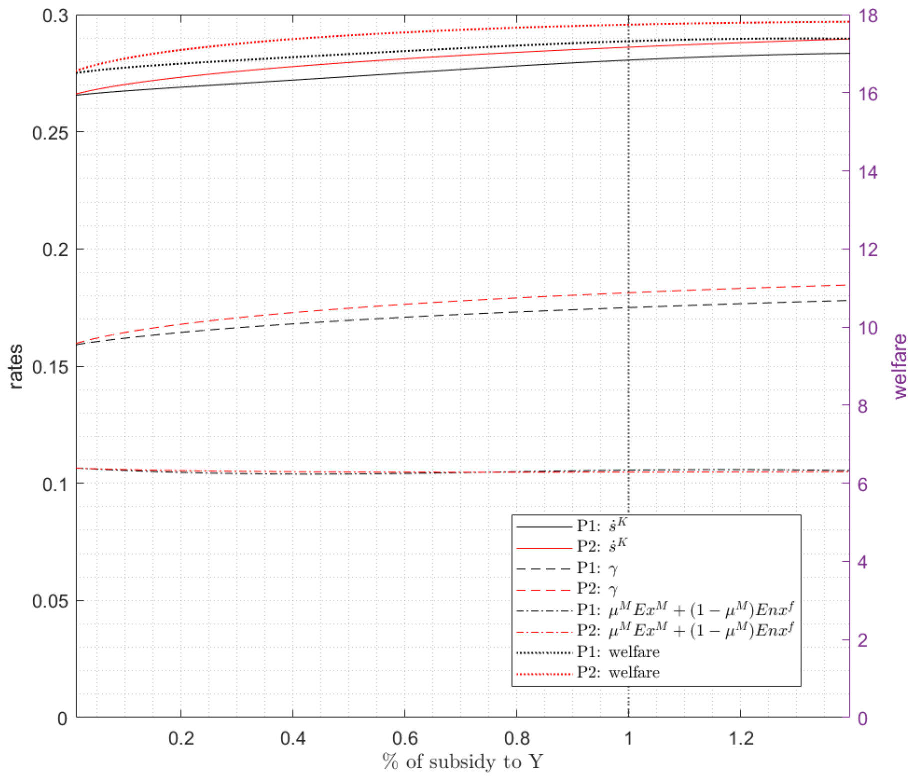

Figure 5 shows comparative statics for both policies across a range of subsidy-to-output ratios. Neither policy substantially changes the average innovation rate of incumbent firms. The marginal returns to R&D are already low in mature, follower-heavy industries, and the subsidies only partially offset the business-stealing effects of entry. Both policies, however, increase the rate of variety creation, with Policy 2 generating a much stronger response. This difference arises because P2 directly raises the payoff from starting a new industry, whereas P1 affects variety creation only indirectly through its effect on firm values.

To compare the two policies at a common fiscal cost, I choose subsidy rates such that government expenditures on each policy amount to 1% of output in the stationary equilibrium. Under Policy 1, this corresponds to rewarding $b_r \approx 0.611$ of the value increase from a successful incumbent innovation. At this level, the long-run growth rate rises by about 5.8 basis points and welfare increases by 5.0%. Under Policy 2, the 1% of output constraint implies a much lower subsidy rate, $b_i \approx 0.087$ of the value of a new monopolist. Even with this smaller subsidy, Policy 2 has larger effects in the calibrated economy: the growth rate increases by about 7.7 basis points and welfare rises by 7.6%.

The life-cycle structure of the model helps explain why the variety-creation instrument is more powerful in this environment. In the calibrated economy, industries follow non-stationary paths in which innovation is front-loaded and incentives weaken as cohorts mature. Subsidizing incumbent R&D yields modest gains because the marginal returns to innovation are already low in the mature, follower-heavy industries that account for a large share of the cross-section. In contrast, subsidizing variety creation directly addresses the entry wedge identified in Section 6.3: it encourages the creation of new, innovation-intensive industries that replenish the cross-section with young cohorts where the marginal returns to both innovation and entry are high. Through this channel, even a modest subsidy to variety creation shifts the cross-sectional distribution toward younger, more dynamic industries and raises aggregate growth more than an equally costly incumbent R&D subsidy in the calibrated economy.

More broadly, these results illustrate the distinctive policy margin that the model creates. In frameworks with a fixed set of industries, innovation policy operates solely on the intensive margin of incumbent R&D. The endogenous industry life-cycle structure in this model opens an additional margin, the rate at which new industries replace mature ones. The quantitative analysis shows that this extensive margin is undervalued in equilibrium and quantitatively important for the transmission of innovation policy.

7 Conclusion

This paper has shown how endogenous industry life-cycles shape aggregate growth and create policy margins that models with a fixed set of industries cannot produce. Industries are born as monopolies, attract followers who compete in patent races, and gradually mature as technology gaps widen and innovation incentives weaken. The cross-sectional distribution of industries over these life-cycle states is stationary in equilibrium and determines the long-run growth rate through an explicit decomposition into variety creation, frontier innovation, and catch-up innovation. The decentralized economy produces too little industry creation relative to the constrained optimum, because private entrants cannot fully appropriate the social return to founding new markets. Equal-cost policy experiments calibrated to U.S. data illustrate the quantitative importance of this margin. Subsidies targeted at variety creation have larger effects on growth and welfare than subsidies to incumbent R&D, because they replenish the cross-section with young, innovation-intensive industries where the marginal returns to both innovation and entry are high.

Two quantitative findings illustrate what the life-cycle structure adds to the analysis of growth and policy. The relationship between startup entry and growth is U-shaped. When entry is low, additional entrants crowd out incumbent innovation and reduce growth, but beyond a threshold, the extensive-margin benefit from industry creation dominates. This non-monotonicity cannot arise in fixed-industry models, where more entry unambiguously intensifies competition. Similarly, tougher product-market competition raises growth in the calibrated economy not by directly increasing incumbent R&D but by redirecting entrants away from mature industries and toward creating new ones. Both results depend on the cross-sectional distribution of industries being an equilibrium object, a feature that is absent when the set of industries is exogenous. The calibrated model also generates non-stationary industry life-cycles, with front-loaded innovation, declining entry, falling prices, and rising output, that are qualitatively consistent with the empirical evidence, even though these patterns are not directly targeted in the calibration.

The framework also provides a tractable way to study changes in business dynamism over time. In the model, higher market-initiation costs reduce variety creation, slow aggregate growth, and shift the industry distribution toward older, less innovative states, while lower initiation costs have the opposite effect. A non-stationary extension in which initiation costs or entrant arrival rates evolve over time could therefore help quantify how changes in entry conditions contributed to the decline in U.S. business dynamism.

A natural next step is to characterize the optimal policy mix over incumbent R&D and industry creation, rather than comparing equal-cost subsidies. Because the model links policy to the endogenous cross-sectional distribution of industries, such an exercise would clarify how the optimal intervention depends on the economy's composition of young and mature industries. More broadly, the analysis suggests that understanding long-run growth requires taking seriously the extensive margin of industry creation, not only the intensive margin of R&D within existing markets.

References

Abernathy, William J, and James M Utterback. 1978. "Patterns of industrial innovation." *Technology review*, 80(7): 40–47.

Agarwal, Rajshree. 1998. "Evolutionary trends of industry variables." *International Journal of Industrial Organization*, 16(4): 511–525.

Agarwal, Rajshree, and Michael Gort. 1996. "The Evolution of Markets and Entry, Exit and Survival of Firms." *The Review of Economics and Statistics*, 78(3): 489–498.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, and Jean-Félix Brouillette. 2026. "Resetting the Innovation Clock: Endogenous Growth through Technological Turnover." Working Paper.

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

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

Akcigit, Ufuk, and Sina T. Ates. 2021. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." *American Economic Journal: Macroeconomics*, 13(1): 257–98.

Akcigit, Ufuk, and Sina T. Ates. 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, Douglas Hanley, and Stefanie Stantcheva. 2022. "Optimal Taxation and R&D Policies." *Econometrica*, 90(2): 645–684.

Argente, David, and Chen Yeh. 2022. "Product Life Cycle, Learning, and Nominal Shocks." *The Review of Economic Studies*, 89(6): 2992–3054.

Argente, David, Doireann Fitzgerald, Sara Moreira, and Anthony Priolo. 2025*a*. "How Do Entrants Build Market Share? The Role of Demand Frictions." *American Economic Review: Insights*, 7(2): 160–76.

Argente, David, Munseob Lee, and Sara Moreira. 2024. "The Life Cycle of Products: Evidence and Implications." *Journal of Political Economy*, 132(2): 337–390.

Argente, David, Salomé Baslandze, Douglas Hanley, and Sara Moreira. 2025*b*. "Patents to Products: Product Innovation and Firm Dynamics." National Bureau of Economic Research Working Paper 34592.

Atkeson, Andrew, and Ariel Burstein. 2019. "Aggregate Implications of Innovation Policy." *Journal of Political Economy*, 127(6): 2625–2683.

Broda, Christian, and David E. Weinstein. 2010. "Product Creation and Destruction: Evidence and Price Implications." *American Economic Review*, 100(3): 691–723.

Cai, Jie, and Can Tian. 2024. "Innovation and Endogenous Knowledge Network Dynamics."

Cavenaile, Laurent, Murat Alp Celik, and Xu Tian. 2025. "Are Markups Too High? Competition, Strategic Innovation, and Industry Dynamics." Working Paper.

Cavenaile, Laurent, Ruben Gaetani, Pau Roldan-Blanco, and Tom Schmitz. 2024. "Industry Life Cycles in General Equilibrium." Rotman School of Management, University of Toronto Working Paper 4892207. SSRN Working Paper.

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

Ferguson, Roger W., and William L. Wascher. 2004. "Distinguished Lecture on Economics in Government: Lessons from Past Productivity Booms." *Journal of Economic Perspectives*, 18(2): 3–28.

Fernald, John. 2014. "A quarterly, utilization-adjusted series on total factor productivity." Federal Reserve Bank of San Francisco.

Filson, Darren. 2001. "The Nature and Effects of Technological Change over the Industry Life Cycle." *Review of Economic Dynamics*, 4(2): 460–494.

Foerster, Andrew, Andreas Hornstein, Pierre-Daniel Sarte, and Mark W Watson. 2025. "The Past and Future of U.S. Structural Change: Compositional Accounting and Forecasting." National Bureau of Economic Research Working Paper 34338.

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

Gort, Michael, and Steven Klepper. 1982. "Time Paths in the Diffusion of Product Innovations." *The Economic Journal*, 92(367): 630–653.

Hobijn, Bart, Martí Mestieri, Nicolas Werquin, and Jing Zhang. 2025. "Quarterly Industry-Level Labor Productivity Data for the U.S." *Economic Perspectives*, , (1): 1–12.

Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. 2008. "A Retrospective Look at the U.S. Productivity Growth Resurgence." *The Journal of Economic Perspectives*, 22(1): 3–24.

Jovanovic, Boyan, and Glenn M. MacDonald. 1994. "The Life Cycle of a Competitive Industry." *Journal of Political Economy*, 102(2): 322–347.

Jovanovic, Boyan, and Zhu Wang. 2025. "Idea Diffusion and Property Rights." Federal Reserve Bank of Richmond Working Paper 20-11R. Revised November 2025.

Klepper, Steven. 1996. "Entry, Exit, Growth, and Innovation over the Product Life Cycle." *The American Economic Review*, 86(3): 562–583.

Liu, Ernest, Atif Mian, and Amir Sufi. 2022. "Low Interest Rates, Market Power, and Productivity Growth." *Econometrica*, 90(1): 193–221.

National Center for Science and Engineering Statistics (NCSES). 2025. "National Patterns of R&D Resources: 2022–23 Data Update." U.S. National Science Foundation NSF 25-326, Alexandria, VA.

Oliner, Stephen D., Daniel E. Sichel, and Kevin J. Stiroh. 2007. "Explaining a Productive Decade." *Brookings Papers on Economic Activity*, 38(1): 81–152.

Olmstead-Rumsey, Jane. 2022. "Market Concentration and the Productivity Slowdown." Working paper, December 24, 2022.

Ribeiro, Bernardo S. C. 2026. "Growth with New and Old Technologies." Manuscript.

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

Sarte, Pierre-Daniel G., and Jack Taylor. 2025. "Sectoral Determinants of Aggregate Productivity Growth." *Economic Brief*, , (25-08).

A Derivations and Proofs

A.1 Lemma 4.1

I define a firm's markup as the ratio of its price to marginal cost, $\eta \equiv p/(w/q)$, where w is the wage and q is the firm's productivity.

Let p_{mn} denote the leader's price in an industry with gap m and n followers and p_{-mn} the common price of a follower, and define the relative price of the leader as $\phi_{mn} \equiv p_{mn}/p_{-mn}$. Given the CES demand system within an industry and the Cobb–Douglas allocation of expenditure across industries, the resulting static pricing game has a unique symmetric Nash equilibrium in prices. The next lemma characterizes the equilibrium markups, quantities, and profits as functions of the state (m, n).

The markup is increasing in the monopolist's lead $\hat{s}$, and the normalized profit $\pi_{\hat{s}}^M$ rises from $1 - \lambda^{-1}$ toward one as the monopolist pulls further ahead of the fringe. The presence of the fringe keeps demand elastic and delivers a well-defined, finite markup for any finite $\hat{s}$.

Lemma A.1 (Pricing decisions)**.** *At time t, consider a monopolist with technology gap $\hat{s}$ over the fringe, born at some date $\tau < t$. Given industry demand $\bar{Y}$ and the presence of the competitive fringe, the monopolist's static profit-maximization problem yields the markup, $\eta_{\hat{s}}^M = \lambda^{\hat{s}}$, the output, $y^M(\tau) = \frac{\bar{Y}}{w}\lambda^{s_\tau^K}$, and the profit, $\Pi_{\hat{s}}^M = \pi_{\hat{s}}^M \bar{Y}$, where $\pi_{\hat{s}}^M = 1 - \lambda^{-\hat{s}}$ is the normalized profit.*

Given technology gap m and n followers, all with productivity λ^{s_f}. Let $\bar{Y}$ denote normalized industry output and w the wage. In the unique symmetric Nash equilibrium in prices, the markups, outputs, and profits of the leader (denoted by subscript mn) and a representative follower (denoted by subscript $-mn$) are:

- Markups*:*

$$\eta_{mn} = \frac{\sigma n + \phi_{mn}^{1-\sigma}}{(\sigma - 1)n},$$

$$\eta_{-mn} = \frac{\sigma\big(\phi_{mn}^{1-\sigma} + n - 1\big) + 1}{(\sigma - 1)\big(\phi_{mn}^{1-\sigma} + n - 1\big)}.$$

- Outputs*:*

$$y_{mn}(s_f) = \frac{(\sigma - 1)n\,\phi_{mn}^{1-\sigma}\,\lambda^{m+s_f}}{\big(\sigma n + \phi_{mn}^{1-\sigma}\big)\big(\phi_{mn}^{1-\sigma} + n\big)}\,\frac{\bar{Y}}{w},$$

$$y_{-mn}(s_f) = \frac{(\sigma - 1)\big(\phi_{mn}^{1-\sigma} + n - 1\big)\lambda^{s_f}}{\big(\sigma(\phi_{mn}^{1-\sigma} + n - 1) + 1\big)\big(\phi_{mn}^{1-\sigma} + n\big)}\,\frac{\bar{Y}}{w}.$$

- Profits*:*

$$\Pi_{mn} = \pi_{mn}\bar{Y}, \qquad \Pi_{-mn} = \pi_{-mn}\bar{Y}, \tag{14}$$

$$\pi_{mn} = \frac{1}{\sigma n \phi_{mn}^{\sigma-1}+1}, \qquad \pi_{-mn} = \frac{1}{\sigma\big(\phi_{mn}^{1-\sigma}+n-1\big)+1}. \tag{15}$$

Moreover, the equilibrium price ratio ϕ_{mn} is uniquely defined by

$$\phi_{mn} n\lambda^m \left(\sigma + \frac{1}{\phi_{mn}^{1-\sigma}+n-1}\right) \;=\; \sigma n + \phi_{mn}^{1-\sigma} \tag{16}$$

for all $m \in \{0,1,\dots\}$ and $n \in \{1,2,\dots\}$.

From the maximization problem, the demand for firm j in industry i is $y_{ji}^d = \frac{p_{ji}^{-\sigma}}{\sum_j p_{ji}^{1-\sigma}}\bar{Y}$. The expressions for markups, outputs, and profits are straightforward from the firm's problem given demand and the price competition. Most of the properties follow from algebra.

1. Following equation can be derived by rearranging the expression.

$$\phi n\lambda^m \left(\sigma + \frac{1}{\phi^{1-\sigma}+n-1}\right) = \sigma n + \phi^{1-\sigma}.$$

   Consider ϕ as a variable for LHS and RHS while other parameters and variables being constant. As LHS is increasing while RHS is decreasing in ϕ, $\lim_{\phi\to 0}\text{LHS} = 0$ while $\lim_{\phi\to 0}\text{RHS} = \infty$, and $\lim_{\phi\to\infty}\text{LHS} = \infty$ while $\lim_{\phi\to\infty}\text{RHS} = \sigma n < \infty$, the solution uniquely exists.

3. Markup changes when m increases and follower markup decreasing in n follow from the previous property. As $\phi_{mn}\lambda^m = \eta_{mn}/\eta_{-mn}$, η_{mn} is also decreasing in n. It means that $\frac{1}{n}\phi_{mn}^{1-\sigma}$ is decreasing in n.

A.2 Lemma 4.2

1. For any normalized value v, $0 < v < \frac{1}{\rho+\gamma}$.
   A normalized profit π a firm can experience satisfies $0 < \pi < 1$ to Lemma 4.1. Therefore, for any states,

$$v < \int_0^\infty e^{-(\rho+\gamma)t}\cdot 1dt = \frac{1}{\rho+\gamma}.$$

   If a firm never invest in R&D, the value is bigger than 0, similarly. Therefore, the sequences are bounded below by 0.

2. $v_{\hat{s}}^M$ is increasing in $\hat{s}$ with $\lim_{\hat{s}\to\infty} v_{\hat{s}}^M = \frac{1+z\theta_M v_{01}}{\rho+\gamma+z\theta_M}$ and $x_{\hat{s}}^M$ is positive and decreasing in $\hat{s}$ with $\lim_{s\to\infty} x_{\hat{s}}^M = 0$.
   For a given s, assume $v_{\hat{s}+1}^M \leq v_{\hat{s}}^M$. Then, $x_{\hat{s}}^M = 0$, so

$$
\begin{aligned}
(\rho+\gamma+z\theta_M)v_{\hat{s}}^M &= \pi_{\hat{s}}^M + z\theta_M v_{01} \\
&< \pi_{\hat{s}+1}^M - R(x_{\hat{s}+1}^M) + x_{\hat{s}+1}^M(v_{\hat{s}+2}^M - v_{\hat{s}+1}^M) + z\theta_M v_{01} \\
&= (\rho+\gamma+z\theta_M)v_{\hat{s}+1}^M.
\end{aligned}
$$

   Sum of the second and third terms in the RHS of the inequality is weakly positive, as the sum can be zero when $x_{\hat{s}+1}^M = 0$. This is an contradiction, so $v_{\hat{s}+1}^M > v_{\hat{s}}^M$ for any s. $\{v_{\hat{s}}^M\}_{s=1}^\infty$ is an increasing and bounded sequence, so a limit exists. From equation (4), $\lim_{s\to} v_{\hat{s}}^M = \frac{1+z\theta_M v_{01}}{\rho+\gamma+z\theta_M}$.
   The positivity of $\{x_{\hat{s}}^M\}_{s=1}^\infty$ comes from increasing $\{v_{\hat{s}}^M\}_{s=1}^\infty$. The normalized value also has a finite limit, so $\lim_{s\to\infty} x_{\hat{s}}^M = 0$.
   Define $\Delta\pi_{\hat{s}}^M = \pi_{\hat{s}+1}^M - \pi_{\hat{s}}^M$, $\Delta v_{\hat{s}}^M = v_{\hat{s}+1}^M - v_{\hat{s}}^M$, and $\Lambda(\Delta v_{\hat{s}}^M) = -R(x_{\hat{s}}^M) + x_{\hat{s}}^M \Delta v_{\hat{s}}^M$. Assume $\Delta v_{\hat{s}}^M > \Delta v_{\hat{s}+1}^M$, but $\Delta v_{\hat{s}-1}^M \leq \Delta v_{\hat{s}}^M$. Then,

$$
\begin{aligned}
(\rho+\gamma+z\theta_M)\Delta v_{\hat{s}}^M &= \Delta\pi_{\hat{s}}^M + \Lambda(\Delta v_{\hat{s}+1}^M) - \Lambda(\Delta v_{\hat{s}}^M) \\
&< \Delta\pi_{\hat{s}-1}^M + \Lambda(\Delta v_{\hat{s}}^M) - \Lambda(\Delta v_{\hat{s}-1}^M) = (\rho+\gamma+z\theta_M)\Delta v_{\hat{s}-1}^M.
\end{aligned}
$$

   The inequality comes from the monotonicity of Λ. Contradiction. Therefore, $\Delta v_{\hat{s}-1}^M > \Delta v_{\hat{s}}^M$. As $\{x_{\hat{s}}^M\}_{s=1}^\infty$ is positive and goes to 0, for any $\hat{s}$, there exists $\hat{s} > s$ such that $x_{\hat{s}}^M > x_{\hat{s}+1}^M$, hence $\Delta v_{\hat{s}}^M > \Delta v_{\hat{s}+1}^M$, so $x_{\hat{s}}^M$ decreases in $\hat{s}$.

3. For any $m \in \{0, 1, \cdots\}$, v_{mn}, v_{-mn}, x_{mn}, and x_{-mn} converge to 0 as $n \to \infty$. For any $n \in \{1, 2, \cdots\}$, v_{mn} converges to $\frac{1}{\rho+\gamma}$ while v_{-mn}, x_{mn}, and x_{-mn} converge to 0 as $m \to \infty$.
   A leader's profit always goes to 1, while a follower's profit always goes to 0 as m goes to infinity, regardless of n. As n goes to infinity, the profits go to 0 for any m. Thus the result follows.

4. A finite maximum number of followers n^* exists.
   As $v_{0n} \to 0$ when $n \to \infty$, there exists $\tilde{n}$ such that $v_{0n} \leq v_1^M - \overline{\kappa}$ for any $n \geq \tilde{n}$. Consider a follower in state mn gets payoff $\pi_{0n'}$, where n' is the current number of followers, until it first becomes a neck-and-neck competitor, while the Markov rates of

the process are the same. Let the value from such process be v^0_{-mn}. Then, if $n \geq \tilde{n}$

$$v_{-mn} \leq v^0_{-mn} \leq \max\left\{\frac{\pi_{0n}}{\rho+\gamma}, v^M_1 - \overline{\kappa}\right\},$$

because any possible follower payoffs after reaching mn is smaller than π_{0n} and for any n, $v_{0n} \leq v^M_1 - \overline{\kappa}$. As π_{0n} also goes to 0 as n increases, there exists some $\hat{n}$ such that $\frac{\pi_{0\hat{n}}}{\rho+\gamma} \leq v^M_1 - \overline{\kappa}$, therefore, $v_{-m\hat{n}} \leq v^M_1 - \overline{\kappa}$ for any m. In other words, there always exists n^* such that, if $n < n^*$, there is at least one m such that $\theta_{mn} > 0$, while $\theta_{mn^*} = 0$ for any m.

A.3 Lemma 4.4

The average levels follow the following differential equations in a stationary equilibrium.

$$\begin{aligned}
\mu_M \dot{\bar{s}}^M_t &= -(z\theta_M + \gamma)\mu_M \bar{s}^M_t + \gamma(s^K_t + 1) + \sum_{\hat{s}}^{\infty} \mu^M_{\hat{s}=1} x^M_{\hat{s}},\\
\mu_M \dot{\bar{s}}^{cf}_t &= -(z\theta_M + \gamma)\mu_M \bar{s}^{cf}_t + \gamma s^K_t,\\
(1-\mu_M)\dot{\bar{s}}^f &= -\gamma(1-\mu_M)\bar{s}^f + z\theta_M \mu_M \bar{s}^M + \sum_{n=1}^{n^*}\sum_{m=1}^{\infty} n x_{-mn}\mu_{mn},\\
(1-\mu_M)\dot{\bar{s}}^l &= -\gamma(1-\mu_M)\bar{s}^l + z\theta_M \mu_M \bar{s}^M + \sum_{n=1}^{n^*}\left((n+1)x_{0n}\mu_0 + \sum_{m=1}^{\infty} x_{mn}\mu_{mn}\right).
\end{aligned}$$

As $\bar{s}^M_t - \bar{s}^{cf}_t = \sum_{\hat{s}=1}^{\infty} \mu^M_{\hat{s}} \hat{s}$ and $\bar{s}^l_t - \bar{s}^f_t = \bar{m} = \sum_{n=1}^{n^*}\sum_{m=0}^{\infty} \mu_{mn} m$ are constant in a stationary equilibrium, one can get the three equations in the lemma using the differential equations, the fact that $\mu_M = \frac{\gamma}{z\theta_M+\gamma}$, and $s^K = \mu_M \bar{s}^M + (1-\mu_M)\bar{s}^f$.

A.4 Proposition 5.1

A monopoly industry with relative level $\hat{s}$ and was born at τ produces $Y^M(\hat{s},\tau) = y^M(\tau) = \frac{\bar{Y}}{w}\lambda^{s^K_\tau}$ until it stops to be a monopolist. Output produced by an oligopoly industry with a technology gap of m and n followers with technology level s^f can be written as $Y_{mn}(s^f) = \tilde{Y}^f_{mn}\lambda^{s^f}\frac{\bar{Y}}{w} = \tilde{Y}^l_{mn}\lambda^{m+s^f}\frac{\bar{Y}}{w}$, where $\tilde{Y}^f_{mn}$ and $\tilde{Y}^l_{mn}$ only depend on m and n. From the labor market conditions, the growth rate of $\frac{\bar{Y}}{w} = \frac{Y}{wN}$ in a stationary equilibrium is $-g_N$.

Over time interval $[t, t+\Delta t]$, changes in industry log outputs are as follows.

$$\Delta \ln Y^M(\hat{s},\tau) = \begin{cases} \hat{s}\ln\lambda + \ln\tilde{Y}^f_{01} = \hat{s}\ln\lambda + \ln\tilde{Y}^l_{01} & \text{with probability } z\theta_M\Delta t + o(\Delta t), \\ -\ln y^M(\tau) & \text{with probability } \delta\Delta t + o(\Delta t), \\ -g_N\Delta t & \text{otherwise}, \end{cases}$$

$$\Delta \ln Y_{0n}(s^f) = \begin{cases} \ln\tilde{Y}^f_{1n} - \ln\tilde{Y}^f_{0n} = \ln\lambda + \ln\tilde{Y}^l_{1n} - \ln\tilde{Y}^l_{0n} & \text{with probability } (n+1)x_{0n}\Delta t + o(\Delta t), \\ \ln\tilde{Y}^f_{0n+1} - \ln\tilde{Y}^f_{0n} = \ln\tilde{Y}^l_{0n+1} - \ln\tilde{Y}^l_{0n} & \text{with probability } z\theta_{0n}\Delta t + o(\Delta t), \\ -\ln Y_{0n}(s^f) & \text{with probability } \delta\Delta t + o(\Delta t), \\ -g_N\Delta t & \text{otherwise}, \end{cases}$$

$$\Delta \ln Y_{mn}(s^f) = \begin{cases} \ln\tilde{Y}^f_{m+1n} - \ln\tilde{Y}^f_{mn} = \ln\lambda + \ln\tilde{Y}^l_{m+1n} - \ln\tilde{Y}^l_{mn} & \text{with probability } x_{mn}\Delta t + o(\Delta t), \\ \ln\lambda + \ln\tilde{Y}^f_{m-1n} - \ln\tilde{Y}^f_{mn} = \ln\tilde{Y}^l_{m-1n} - \ln\tilde{Y}^l_{mn} & \text{with probability } nx_{-mn}\Delta t + o(\Delta t), \\ \ln\tilde{Y}^f_{mn+1} - \ln\tilde{Y}^f_{mn} = \ln\tilde{Y}^l_{mn+1} - \ln\tilde{Y}^l_{mn} & \text{with probability } z\theta_{mn}\Delta t + o(\Delta t), \\ -\ln Y_{mn}(s^f) & \text{with probability } \delta\Delta t + o(\Delta t), \\ -g_N\Delta t & \text{otherwise}. \end{cases}$$

In cases except for destruction of industries, outputs also decrease by $g_N\Delta t$ in addition to the changes from state transition, but it is ignored as it would not appear in the limit. Also, at each moment t, a new industry with output $Y^M(1,t)$ appears with rate γ per each industry. Therefore,

$$\frac{d}{dt}\int_0^{N_t} \ln Y_{it}di = \gamma N_t \ln Y^M(1,t) + \mu_M N_t \mathbb{E}_t \frac{d}{dt}\ln Y^M + (1-\mu_M)N_t\mathbb{E}_t\frac{d}{dt}\ln Y_{mn}.$$

Let F^f_{mnt} be a distribution of followers' technology levels for the industries with gap m and n followers at time t and F^M_{cft} be a distribution of fringe firms' technology levels for the monopolies at t. Using the stationary distribution of technology gaps, the aggregate growth rate is as follows.

$$
\begin{aligned}
g_t &= g_N + \frac{d}{dt}\ln \bar{Y}_t = g_N - \frac{\dot{N}_t}{N_t^2}\int_0^{N_t}\ln Y_{it}di + \frac{1}{N_t}\frac{d}{dt}\int_0^{N_t}\ln Y_{it}di \\
&= g_N - g_N\ln\bar{Y}_t + \gamma\ln Y^M(1,t) - \delta\ln\bar{Y}_t - g_N + \sum_{\hat{s}=1}^{\infty}\mu_{\hat{s}}^M z\theta_M(\hat{s}\ln\lambda + \ln\tilde{Y}_{01}^f) \\
&\quad + \ln\lambda\sum_{n=1}^{n^*}\sum_{m=1}^{\infty}nx_{-mn}\mu_{mn} + \sum_{n=1}^{n^*}\Bigg(\mu_{0n}\Big[(n+1)x_{0n}(\ln\tilde{Y}_{1n}^f - \ln\tilde{Y}_{0n}^f) + z\theta_{0n}(\ln\tilde{Y}_{0n+1}^f - \ln\tilde{Y}_{0n}^f)\Big] \\
&\quad + \sum_{m=1}^{\infty}\mu_{mn}\Big[x_{mn}(\ln\tilde{Y}_{m+1n}^f - \ln\tilde{Y}_{mn}^f) + nx_{-mn}(\ln\lambda + \ln\tilde{Y}_{m-1n}^f - \ln\tilde{Y}_{mn}^f) + z\theta_{mn}(\ln\tilde{Y}_{mn+1}^f - \ln\tilde{Y}_{mn}^f)\Big]\Bigg) \\
&= \; - g_N\ln\bar{Y}_t + \gamma(\ln\frac{\bar{Y}_t}{w_t} + s_t^K\ln\lambda) + (g_N - \gamma)\ln\bar{Y}_t + z\theta_M\ln\lambda\sum_{\hat{s}=1}^{\infty}\mu_{\hat{s}}^M\hat{s} + \ln\lambda\sum_{n=1}^{n^*}\sum_{m=1}^{\infty}nx_{-mn}\mu_{mn} \\
&\quad + \ln\tilde{Y}_{01}^f(x_{-11}\mu_{11} - 2x_{01}\mu_{01} + z\theta_M\mu_M - z\theta_{01}\mu_{01}) + \cdots \\
&\quad + \ln\tilde{Y}_{mn}^f(x_{m-1n}\mu_{m-1n} + nx_{-m-1n}\mu_{m+1n} - (x_{mn} + nx_{-mn})\mu_{mn} + z\theta_{mn-1}\mu_{mn-1} - z\theta_{mn}\mu_{mn}) \\
&\quad + \cdots
\end{aligned}
$$

$$
\begin{aligned}
&= \gamma s_t^K\ln\lambda - \gamma\mu_M\ln\lambda\int_0^{\infty}sdF_{cft}^M(s) - \gamma\sum_{n=1}^{n^*}\sum_{m=0}^{\infty}\mu_{mn}\int_1^{\infty}\left(\ln\tilde{Y}_{mn}^f + s\ln\lambda\right)dF_{mnt}^f(s) \\
&\quad + z\theta_M\mu_M\bar{s}^M\ln\lambda - z\theta_M\mu_M\ln\lambda\int_0^{\infty}sdF_{cft}^M(s) + \ln\lambda\sum_{n=1}^{n^*}\sum_{m=1}^{\infty}nx_{-mn}\mu_{mn} + \gamma\sum_{n=1}^{n^*}\sum_{m=0}^{\infty}\mu_{mn}\ln\tilde{Y}_{mn}^f \\
&= \; \ln\lambda\left[\gamma s_t^K + z\theta_M\mu_M\left(\bar{s}_t^M - \bar{s}_t^{cf}\right) - \gamma\left(\mu_M\bar{s}_t^{cf} + (1-\mu_M)\bar{s}_t^f\right) + \sum_{n=1}^{n^*}\sum_{m=1}^{\infty}nx_{-mn}\mu_{mn}\right]
\end{aligned}
$$

Using $s_t^K = \mu_M\bar{s}_t^M + (1-\mu_M)\bar{s}_t^f$ and $\mu_M = \frac{\gamma}{\gamma+z\theta_M}$ in a stationary equilibrium, it can be shown that the sum of the first three terms in the bracket is equal to $\gamma\left(\bar{s}_t^M - \bar{s}_t^{cf}\right)$. From Lemma 4.4,

$$
g = \ln\lambda\left(\gamma + \sum_{\hat{s}=1}^{\infty}\mu_{\hat{s}}^M x_{\hat{s}}^M + \sum_{n=1}^{n^*}\sum_{m=1}^{\infty}nx_{-mn}\mu_{mn}\right) = \ln\lambda\cdot\dot{s}^K.
$$

Similarly,

$$
g = \ln\lambda\left(\gamma + \sum_{\hat{s}=1}^{\infty}\mu_{\hat{s}}^M x_{\hat{s}}^M + \sum_{n=1}^{n^*}\left((n+1)x_{0n}\mu_0 + \sum_{m=1}^{\infty}x_{mn}\mu_{mn}\right) - \gamma(1-\mu_M)\bar{m}\right).
$$

A.5 Proof of Proposition 5.2

For part (i), fix an oligopoly state (m, n). Over a short interval Δ, the gap changes only through innovation events that either widen or narrow the leader's advantage.

If $m = 0$, the gap can only widen, and does so by one with probability

$$(n+1)x_{0n}\Delta + o(\Delta).$$

Hence

$$\lim_{\Delta \downarrow 0} \frac{1}{\Delta}\mathbb{E}[\tilde{m}_\Delta - \tilde{m} \mid \tilde{m} = 0, \tilde{n} = n] = (n+1)x_{0n}.$$

If $m \geq 1$, leader innovation widens the gap by one with probability

$$x_{mn}\Delta + o(\Delta),$$

while follower catch-up narrows it by one with probability

$$nx_{-mn}\Delta + o(\Delta).$$

Therefore

$$\lim_{\Delta \downarrow 0} \frac{1}{\Delta}\mathbb{E}[\tilde{m}_\Delta - \tilde{m} \mid \tilde{m} = m, \tilde{n} = n] = x_{mn} - nx_{-mn}.$$

Taking expectations over the stationary cross-sectional distribution conditional on being an oligopoly gives

$$\lim_{\Delta \downarrow 0} \frac{1}{\Delta}\mathbb{E}[\tilde{m}_\Delta - \tilde{m}] = \frac{1}{1-\mu_M}\left[\sum_{n=1}^{n^*}(n+1)x_{0n}\mu_{0n} + \sum_{n=1}^{n^*}\sum_{m=1}^{\infty} x_{mn}\mu_{mn} - \sum_{n=1}^{n^*}\sum_{m=1}^{\infty} nx_{-mn}\mu_{mn}\right].$$

By Lemma 4.4, the bracketed term equals $\gamma(1-\mu_M)\bar{m}$. Hence

$$\lim_{\Delta \downarrow 0} \frac{1}{\Delta}\mathbb{E}[\tilde{m}_\Delta - \tilde{m}] = \gamma\bar{m}.$$

Since $\bar{m} > 0$, this proves part (i).

For part (ii), recall that $U \sim \text{Exp}(\gamma)$ denotes the age of a randomly sampled surviving oligopoly from the stationary cross-section. By definition,

$$\bar{m} = \mathbb{E}[f(U)].$$

Fix $\varepsilon > 0$. By the law of iterated expectations,

$$\bar{m} = \mathbb{P}(U \le \varepsilon)\,\mathbb{E}[f(U) \mid U \le \varepsilon] + \mathbb{P}(U > \varepsilon)\,\mathbb{E}[f(U) \mid U > \varepsilon].$$

Using $\bar{m}_\varepsilon = \mathbb{E}[f(U) \mid U > \varepsilon]$, this becomes

$$\bar{m} = (1 - e^{-\gamma\varepsilon})\,\mathbb{E}[f(U) \mid U \le \varepsilon] + e^{-\gamma\varepsilon}\bar{m}_\varepsilon.$$

Now $f(0) = 0$, since a newborn oligopoly starts at zero gap. Moreover, $f(a) \to 0$ as $a \downarrow 0$, so

$$\mathbb{E}[f(U) \mid U \le \varepsilon] \to 0 \qquad \text{as } \varepsilon \downarrow 0.$$

Because $\bar{m} > 0$, there exists $\varepsilon_0 > 0$ such that

$$\mathbb{E}[f(U) \mid U \le \varepsilon] < \bar{m} \qquad \text{for all } \varepsilon \in (0, \varepsilon_0].$$

Rearranging the mixture identity then yields

$$\bar{m}_\varepsilon = \frac{\bar{m} - (1 - e^{-\gamma\varepsilon})\mathbb{E}[f(U) \mid U \le \varepsilon]}{e^{-\gamma\varepsilon}} > \bar{m} \qquad \text{for all } \varepsilon \in (0, \varepsilon_0].$$

This proves part (ii).

A.6 Proof of Corollary ??

Because industry exit occurs at the exogenous state-independent rate δ, conditioning on survival to age a does not distort the within-oligopoly transition dynamics. Hence the derivative of the cohort mean gap is given by the expected instantaneous change in the gap conditional on survival:

$$f'(a) = \lim_{\Delta \downarrow 0} \frac{1}{\Delta}\mathbb{E}[\hat{m}_{a+\Delta} - \hat{m}_a \mid \text{alive at age } a]\,.$$

Under Assumption 5.1, this quantity is weakly positive for all sufficiently large ages. Therefore

$$f'(a) \ge 0$$

for all sufficiently large a, which implies that $f(a)$ is weakly increasing for sufficiently large ages.

A.7 Proof of Proposition 5.3

For part (i), conditional on remaining a monopoly and surviving, the monopolist's quality changes only when it innovates. Hence

$$\lim_{\Delta \downarrow 0} \frac{1}{\Delta} \mathbb{E}\big[\ln q^M_{a+\Delta} - \ln q^M_a \,\big|\, \text{still monopoly, alive}\big] = \ln \lambda \, \mathbb{E}\big[x^M_{\hat{s}_a} \,\big|\, \text{still monopoly, alive}\big] .$$

While the industry remains a monopoly, the lead $\hat{s}_a$ is weakly increasing along every history, since monopoly innovation can only increase the lead. By Lemma 4.2, $x^M_{\hat{s}}$ is decreasing in $\hat{s}$ and satisfies $\lim_{\hat{s}\to\infty} x^M_{\hat{s}} = 0$. Therefore

$$\mathbb{E}\big[x^M_{\hat{s}_a} \,\big|\, \text{still monopoly, alive}\big] \searrow 0 \qquad \text{as } a \to \infty,$$

which proves part (i).

For part (ii), write the oligopoly industry price as

$$P_{mn}(s^f) = \frac{w}{\lambda^{s^f}} \Phi(m,n),$$

where $\Phi(m,n)$ depends only on the relative state (m,n). Using the equilibrium price ratio ϕ_{mn} from Lemma **??**, define

$$A_{mn} \equiv \phi_{mn}^{1-\sigma}.$$

Then the oligopoly price factor can be written as

$$\Phi(m,n) = \frac{\sigma(A_{mn} + n - 1) + 1}{(\sigma - 1)(A_{mn} + n - 1)} \left(A_{mn} + n\right)^{\frac{1}{1-\sigma}} .$$

Moreover, the equilibrium condition for ϕ_{mn} implies

$$n\phi_{mn}\lambda^m \left(\sigma + \frac{1}{\phi_{mn}^{1-\sigma} + n - 1}\right) = \sigma n + \phi_{mn}^{1-\sigma}.$$

Fix n. As $m \to \infty$, the previous equation implies

$$\phi_{mn} \to 0, \qquad \phi_{mn} \sim (\sigma n)^{-1/\sigma} \lambda^{-m/\sigma}.$$

Substituting into the expression for $\Phi(m,n)$ yields

$$\Phi(m,n) \sim \frac{\sigma}{\sigma - 1} (\sigma n)^{-1/\sigma} \lambda^{-m/\sigma}.$$

Hence, for each fixed $n \in \{1, \dots, n^*\}$, there exists $M_P(n) < \infty$ such that

$$\Phi(m+1, n) \leq \Phi(m, n) \qquad \text{for all } m \geq M_P(n).$$

Since $P_{mn}(s^f)/w = \lambda^{-s^f}\Phi(m, n)$, this implies

$$\frac{P_{m+1,n}(s^f)}{w} \leq \frac{P_{mn}(s^f)}{w} \qquad \text{for all } m \geq M_P(n).$$

For the dependence on n, Lemma 4.4 implies that, for each fixed m, the wage-adjusted industry price is weakly decreasing in n. Therefore, for each $n \in \{1, \dots, n^* - 1\}$,

$$\frac{P_{m,n+1}(s^f)}{w} \leq \frac{P_{mn}(s^f)}{w}.$$

Finally, each active industry earns the same revenue flow $\bar{Y}$, so

$$Y_{mn}(s^f) = \frac{\bar{Y}}{P_{mn}(s^f)}.$$

Therefore the previous price inequalities imply

$$Y_{m+1,n}(s^f) \geq Y_{mn}(s^f) \qquad \text{for all } m \geq M_P(n),$$

and

$$Y_{m,n+1}(s^f) \geq Y_{mn}(s^f) \qquad \text{for each } m \text{ and } n \in \{1, \dots, n^* - 1\}.$$

This proves part (ii).

For part (iii), Corollary **??** implies that, under Assumption 5.1, the expected technology gap of surviving oligopoly cohorts is weakly increasing for sufficiently large ages. In addition, the number of followers cannot decrease along the industry life-cycle. Hence sufficiently old surviving oligopoly cohorts move toward states with larger gaps and weakly larger follower counts.

By part (ii), larger gaps and larger follower counts imply lower relative prices and higher quantities in sufficiently mature oligopoly states. By Lemma 4.2, for fixed n,

$$\lim_{m \to \infty} x_{mn} = \lim_{m \to \infty} x_{-mn} = 0,$$

and entry becomes unprofitable once follower counts are sufficiently high. Thus, as surviving oligopoly cohorts become more mature, innovation incentives weaken and entry becomes less attractive. Finally, leadership turnover within any fixed unit interval becomes less likely in large-gap states, because follower catch-up probabilities become small and a turnover event requires the gap to be reduced before leadership can change. Therefore, for sufficiently large ages, surviving oligopoly cohorts move toward lower-price, higher-quantity, lower-innovation, lower-entry, and lower-turnover states. This proves part (iii).

B Equilibrium Definition

B.1 Hamilton-Jacobi-Bellman Equations and Normalization

The Hamilton-Jacobi-Bellman equations are as follows.

$$
\begin{aligned}
rV_{\hat{s}}^M - \dot{V}_{\hat{s}}^M &= \max_{x\in[0,\infty)} \Pi_{\hat{s}}^M - R(x)\bar{Y} + x(V_{\hat{s}+1}^M - V_{\hat{s}}^M) + z\theta_M(V_{01} - V_{\hat{s}}^M) - \delta V_{\hat{s}}^M \quad \text{for } s \in \{1, 2, \cdots\}, \\
rV_{mn} - \dot{V}_{mn} &= \max_{x\in[0,\infty)} \Pi_{mn} - R(x)\bar{Y} + x(V_{m+1n} - V_{mn}) + nx_{-mn}(V_{m-1n} - V_{mn}) \\
&\qquad + z\theta_{mn}(V_{mn+1} - V_{mn}) - \delta V_{mn} \quad \text{for } m, n \in \{1, 2, \cdots\}, \\
rV_{-mn} - \dot{V}_{-mn} &= \max_{x\in[0,\infty)} \Pi_{-mn} - R(x)\bar{Y} + (x + (n-1)x_{-mn})(V_{-m+1n} - V_{-mn}) + x_{mn}(V_{-m-1n} - V_{-mn}) \\
&\qquad + z\theta_{mn}(V_{-mn+1} - V_{-mn}) - \delta V_{-mn} \quad \text{for } m, n \in \{1, 2, \cdots\}, \\
rV_{0n} - \dot{V}_{0n} &= \max_{x\in[0,\infty)} \Pi_{0n} - R(x)\bar{Y} + x(V_{1n} - V_{0n}) + nx_{0n}(V_{-1n} - V_{0n}) \\
&\qquad + z\theta_{0n}(V_{0n+1} - V_{0n}) - \delta V_{0n} \quad \text{for } n \in \{1, 2, \cdots\},
\end{aligned}
$$

where θs are the weighted entry decision as previously defined.

Define $v := \frac{V}{\bar{Y}}$, with corresponding state sub/superscripts. As $\dot{V} = \dot{v}\bar{Y} + v\dot{\bar{Y}}$ and $\frac{\dot{\bar{Y}}}{\bar{Y}} = g - g_N = r - \rho - g_N$,

$$rV - \dot{V} = rv\bar{Y} - \dot{v}\bar{Y} - v\dot{\bar{Y}} = (\rho + g_N)v\bar{Y} - \dot{v}\bar{Y}.$$

In a stationary equilibrium, the RHS of normalized HJBs become time-invariant, so $\dot{v} = 0$ for any states.

B.2 Kolmogorov forward equations

The Kolmogorov forward equations for industry shares are as follows. In a stationary equilibrium, $\dot{\mu} = 0$.

$$(\delta + g_N)\mu_1^M + \dot{\mu}_1^M = -x_1^M \mu_1^M + \gamma - z\theta_M \mu_1^M, \tag{17}$$

$$(\delta + g_N)\mu_{\hat{s}}^M + \dot{\mu}_{\hat{s}}^M = x_{\hat{s}-1}^M \mu_{\hat{s}-1}^M - x_{\hat{s}}^M \mu_{\hat{s}}^M - z\theta_M \mu_{\hat{s}}^M \quad \text{for } s \geq 2, \tag{18}$$

$$(\delta + g_N)\mu_{01} + \dot{\mu}_{01} = x_{-11}\mu_{11} - 2x_{01}\mu_{01} + z\theta_M \mu_M - z\theta_{01}\mu_{01}, \tag{19}$$

$$(\delta + g_N)\mu_{0n} + \dot{\mu_{0n}} = nx_{-1n}\mu_{1n} - (n+1)x_{0n}\mu_{0n} + z\theta_{0n-1}\mu_{0n-1} - z\theta_{0n}\mu_{0n} \quad \text{for } n \geq 2, \tag{20}$$

$$(\delta + g_N)\mu_{11} + \dot{\mu_{11}} = 2x_{01}\mu_{01} + x_{-21}\mu_{21} - (x_{11} + x_{-11})\mu_{11} - z\theta_{11}\mu_{11}, \tag{21}$$

$$\begin{aligned}(\delta + g_N)\mu_{1n} + \dot{\mu_{1n}} = &(n+1)x_{0n}\mu_{0n} + nx_{-2n}\mu_{2n} - (x_{1n} + nx_{-1n})\mu_{1n} \\ &+ z\theta_{1n-1}\mu_{1n-1} - z\theta_{1n}\mu_{1n} \quad \text{for } n \geq 2,\end{aligned} \tag{22}$$

$$\begin{aligned}(\delta + g_N)\mu_{mn} + \dot{\mu_{mn}} = &x_{m-1n}\mu_{m-1n} + nx_{-m-1n}\mu_{m+1n} - (x_{mn} + nx_{-mn})\mu_{mn} \\ &+ z\theta_{mn-1}\mu_{mn-1} - z\theta_{mn}\mu_{mn} \quad \text{for } m \geq 1 \text{ and } n \geq 2,\end{aligned} \tag{23}$$

B.3 Market clearing conditions

$$\begin{aligned}
Y &= C + Y\left(\mathcal{R} + \mathcal{K}\right), \\
&\quad \text{where } \mathcal{R} = \sum_{\hat{s}=1}^{\infty} \mu_{\hat{s}}^M R(x_{\hat{s}}^M) + \sum_{m=0}^{\infty}\sum_{n=1}^{\infty} \mu_{mn}\left(R(x_{mn}) + nR(x_{-mn})\right) \\
&\quad \text{and } \mathcal{K} = z\left(\mu_M \int_{\underline{\kappa}}^{\hat{\kappa}_M} \kappa dG(\kappa) + \sum_{m=0}^{\infty}\sum_{n=1}^{\infty} \mu_{mn} \int_{\underline{\kappa}}^{\hat{\kappa}_{mn}} \kappa dG(\kappa)\right), \\
1 &= \frac{Y}{w}\left(\sum_{\hat{s}=1}^{\infty} \mu_{\hat{s}}^M \lambda^{-\hat{s}} + \sum_{m=0}^{\infty}\sum_{n=1}^{\infty} \mu_{mn}\left(\frac{n(\sigma-1)\phi_{mn}^{1-\sigma}}{(\sigma n + \phi_{mn}^{1-\sigma})(\phi_{mn}^{1-\sigma} + n)}\right.\right. \\
&\qquad\qquad \left.\left. + \frac{n(\sigma-1)(\phi_{mn}^{1-\sigma} + n - 1)}{(\sigma(\phi_{mn}^{1-\sigma} + n - 1) + 1)(\phi_m^{1-\sigma} + n)}\right)\right). \\
A &= Y\left(\sum_{\hat{s}=1}^{\infty} \mu_{\hat{s}}^M v_{\hat{s}}^M + \sum_{m=0}^{\infty}\sum_{n=1}^{\infty} \mu_M (v_m + v_{-m})\right).
\end{aligned}$$

C Quantitative Appendix

C.1 Computation Procedure

For computation, I need to set a maximum level a monopolist can reach and a maximum gap and a maximum number of followers that an oligopoly can achieve. The numbers I choose are

30, 50, and 30 in order which are large enough for calculating the outcomes. Main algorithm of the quantitative model is as follows.

1. Guess γ and v_{01}.

2. Solve for monopoly values.

3. Solve for oligopoly values using the new v_1^M.

4. Solve for stationary distribution and new γ.

5. Check if the new γ and v_{01} match the initial guess. If not, go back to the first step.

C.2 Transition Probability and Industry Variables

From the process of industry states, I can define a transition probability between different states after t amount of time, conditional on the industry's survival, as follows.

$$P_{x_1 x_2}(t) = P(state(t+h) = x_2 | state(h) = x_1 \text{ and not destroyed until } t+h)$$

where $state(t) \in \{1, 2, \cdots, s, \cdots\} \cup \{01, 02, \cdots, mn, \cdots\}$ is the state of the firm at time t. The corresponding Kolmogorov backward and forward equations are as follows.

$$\dot{P}(t) = RP(t)$$
$$\dot{P}(t) = P(t)R$$

where R is a transition rate matrix of the process conditional of the survival of the industry. With the Kolmogorov equations and R, the transition probabilities conditional on survival can be derived as $P(t) = e^{Rt}$.

C.3 Discontinuities in Comparative Statics

The comparative statics in Figures 3 and 4 exhibit several discrete jumps in growth and innovation outcomes as parameters vary. These discontinuities arise from the discrete industry state space and the bounded support of market initiation costs, not from numerical instability.

Entry decisions are governed by threshold comparisons among value functions: for a given distribution of costs, a potential entrant compares the value of joining an incumbent industry in state S with the value of creating a new industry. As parameters change, the ordering of values such as v_{02} and v_1^M can switch. Because G has bounded support, the economy can hit a corner in which the entry condition for a given state moves from interior to corner, shifting the probability of entry in that state from strictly positive to zero (or vice versa) for the entire cost distribution. These nonlinear switches in state-contingent entry generate abrupt changes in the variety creation rate γ. Since γ enters the effective discount rate $(\rho + g_N)$, these shifts feed back into firm values, generating the discrete changes in aggregate rates and welfare visible in the figures. This feature reflects the inherently discrete and non-linear nature of industrial evolution in the model.