Resetting the Innovation Clock: Endogenous Growth through Technological Turnover*

Philippe Aghion[1] Antonin Bergeaud[2]

Timo Boppart[3] Jean-Félix Brouillette[4]

January 22, 2026

Abstract

We propose a model of endogenous economic growth with "weak" scale effects and diminishing returns to innovation at the micro level. In our model, entrants introduce new technologies through *research* and incumbents incrementally improve them through *development*. Over time, further improvement becomes harder such that firms ultimately run out of ideas and exit, paving the way for entrants that discover new technologies with further room for improvement. This turnover gives rise to a continuous stream of (temporary) opportunities for technological improvements that sustain economic growth. In a stationary equilibrium, the long-run growth rate is constant and *endogenous* to market incentives.

*We are grateful to Conor Walsh, Phillip Trammell, Chad Jones, and Pete Klenow for helpful comments. Boppart thanks SNSF consolidator grant TMCG-1_223114 for financial support.

[1]Collège de France, INSEAD, LSE, and CEPR. E-mail: philippe.aghion@insead.edu

[2]HEC Paris, CEP-LSE, and CEPR. E-mail: bergeaud@hec.fr

[3]IIES, Stockholm University, and University of Zurich. E-mail: timo.boppart@iies.su.se

[4]HEC Montréal. E-mail: jean-felix.brouillette@hec.ca

1 Introduction

In this paper, we propose a model of innovation-driven growth that jointly replicates two empirical regularities that prior endogenous-growth frameworks have struggled to capture. First, economic growth has not accelerated despite sustained population growth (Jones, 1995). Second, within any specific technological field, discovering better ideas becomes increasingly difficult (Bloom, Jones, Van Reenen and Webb, 2020). However, even as new ideas become harder to find at both the macro and micro levels, our model delivers a constant rate of long-run economic growth that is endogenous to market incentives, particularly tax and R&D policy (Akcigit, Grigsby, Nicholas and Stantcheva, 2021; Dechezleprêtre, Einiö, Martin, Nguyen and Van Reenen, 2023).

To achieve this, our model distinguishes between *research* and *development*, reminiscent of Aghion and Howitt (1996). Research is conducted by entrants and delivers fundamentally new products—what Mokyr (1992) calls "macro inventions." Development is carried out by incumbents and consists of improvements to existing products through incremental "micro inventions." This distinction enables our model to reconcile endogenous growth with evidence that ideas are becoming harder to find. Development faces diminishing returns within a product line, but research continually introduces new products with fresh development opportunities, effectively "resetting the innovation clock." As these new products arrive, they erode incumbents' market shares and eventually displace mature technologies. The resulting cycle of research, development, and creative destruction delivers a continuous stream of *temporary* innovation opportunities that sustain long-run growth.

The evolution of the camera industry illustrates this distinction clearly. Early film cameras were *developed* from simple box models into sophisticated single-lens reflex systems. However, as improvement opportunities in film-based imaging were exhausted, *research* delivered a fundamentally new product—the digital camera—which opened fresh avenues for subsequent *development* in sensor technology, image processing, and software-driven features. Mokyr (1990) documents similar patterns in long-distance communications, where the telegraph, long-wave radio, and shortwave radio each represented macro inventions that triggered a subsequent stream of micro inventions.[1]

Two crucial but empirically motivated assumptions underlie our framework. First, new entrants "stand on the shoulders of giants": the initial quality of a new product scales with the contemporaneous average quality across products. This assumption is necessary to sustain quality-led growth; without it, average quality growth would stall

[1]Similar patterns appear in lighting (gas → incandescent → LED), sound recording (phonograph → magnetic tape → digital), data storage (magnetic tape → HDD → SSD), and semiconductors (vacuum tubes → transistors → integrated circuits).

due to the inevitable exhaustion of development opportunities within product lines, and our model would revert to a semi-endogenous framework. Second, we assume that entry costs scale with aggregate productivity. Combined with free entry into research, as in second-generation endogenous growth models (Peretto, 1998; Dinopoulos and Thompson, 1998; Young, 1998; Segerstrom, 1998; Howitt, 1999), this rules out "strong" scale effects by ensuring that the number of product lines scales with the population and the ratio of developers per line remains constant.[2] We elaborate on these assumptions and the supporting evidence in Section 3.

With these assumptions, our model admits a tractable balanced growth path where the growth rate of output per capita is constant and, crucially, *endogenous* to policy. Growth is fueled by two distinct forces: the expansion of product variety and product quality improvements. The former is tied to population growth and is unresponsive to policy, as in semi-endogenous growth models. The latter is driven by development incentives and is responsive to policy (e.g., research and development subsidies), as in endogenous growth models. Importantly, if population growth falls to zero, our model still delivers positive endogenous growth due to quality upgrading, which is insulated from demographic forces.

To assess the quantitative implications of our framework, we calibrate the model to U.S. establishment-level data, treating establishments as proxies for products following Garcia-Macia, Hsieh and Klenow (2019). Our deliberately parsimonious model replicates several moments from the U.S. Business Dynamics Statistics, including the exit rate, average establishment size, and the shape of the size distribution. This simple calibration exercise allows us to quantify counterfactuals and the impact of different policies. For instance, we find that a 10% subsidy to development expenses increases the long-run growth rate by 50 basis points, while a population growth slowdown from 1% to zero reduces long-run per-capita growth from 2% to 1.78%.

Literature review

This paper makes two contributions. First, it speaks to the long-standing discussion on the future of technological progress, contested by techno-pessimists like Gordon (2017), who argue that transformative innovations are behind us, and optimists like Mokyr (2014), who see recent technological revolutions (e.g., IT and more recently AI) as sources of potentially unprecedented future growth. Our model offers a synthesis: it formally incorporates the view that ideas become harder to find within any single

[2] Another way to sterilize this strong scale effect is to assume that vertical innovation faces "severely" diminishing returns to development labor as in Trammell (2025).

technological paradigm, yet it shows how the continual discovery of new paradigms can sustain aggregate economic growth indefinitely.

Second, our work contributes to the debate between *endogenous* and *semi-endogenous* growth paradigms. This debate has persisted since the mid-1990s, largely because of a seemingly insurmountable empirical challenge; convincingly discriminating whether market incentives (e.g., market size or R&D policies) have a causal effect on the *long-run* level or growth rate of aggregate productivity when transition dynamics can be slow. Consequently, the discussion has advanced by developing theoretical models from both sides and testing them against other, more measurable empirical moments.

The first generation of endogenous growth models (e.g., Romer (1990) and Aghion and Howitt (1992)) developed frameworks in which market incentives could shape long-run economic growth. However, in these models, the population had to remain constant to prevent explosive growth, which ran counter to empirical evidence. In response, the semi-endogenous growth model of Jones (1995) eliminated this "strong scale effect" by assuming that innovation efficiency declines with the aggregate stock of knowledge. This solution, however, came at the cost of rendering long-run productivity growth unresponsive to market incentives and having the strong prediction that long-run growth is impossible without population growth.

A "second generation" of models subsequently sought to break the strong scale effect while preserving a role for market incentives (e.g., Peretto (1998), Dinopoulos and Thompson (1998), Young (1998), Segerstrom (1998), and Howitt (1999)). This was achieved by allowing the number of firms to grow in proportion to the population (through free entry), keeping the number of developers per firm constant and delivering constant growth at the firm level. However, recent evidence from Bloom et al. (2020) directly challenges the assumption that a constant number of developers per firm can deliver constant firm-level growth. They show across multiple contexts that micro-level innovation efficiency is instead consistently declining.

Our contribution to this debate is to propose a model that reconciles this recent empirical evidence with the desirable predictions of the second wave of endogenous growth models: long-run economic growth is responsive to market incentives, yet does not explode even as the population expands. In our model, this reconciliation is achieved by the endogenous arrival of technological breakthroughs that continuously reset the innovation clock. In our framework, these breakthroughs are conducted exclusively by entrants, while incumbents focus on incremental development. This is a simplification; in reality, incumbents also introduce breakthroughs (e.g., Kodak developed the first digital camera).

2 A new growth model

In this section, we present an intentionally simple model that captures the essence of our argument: firms invest in the *development* of existing products to improve their quality, but eventually run out of ideas and exit. Simultaneously, investments in *research* lead to the entry of new firms with entirely new product lines. As a consequence, quality-upgrading opportunities are continuously replenished in aggregate even though they are ultimately exhausted on any existing line.

2.1 The economic environment

Consider a continuous-time economy where time is indexed by $t \in [0, \infty)$. This economy is populated by a representative household of measure N_t evolving according to:

$$\dot{N}_t = n \cdot N_t, \tag{1}$$

where $n \geq 0$. The household inelastically supplies one unit of labor at every point in time, and has logarithmic preferences over per-capita consumption c_t such that lifetime utility is defined as:

$$U_0 = \int_0^\infty e^{-(\rho - n)t} \ln(c_t) \mathrm{d}t. \tag{2}$$

Here, $\rho > n$ denotes the rate of time preference. The individual consumption basket is a Dixit and Stiglitz (1977) aggregator of differentiated products indexed by $i \in \mathcal{I}_t$:

$$c_t = \left(\int_{i \in \mathcal{I}_t} (q_{it} \cdot c_{it})^{\frac{\theta - 1}{\theta}} \mathrm{d}i \right)^{\frac{\theta}{\theta - 1}}, \tag{3}$$

where c_{it} is the consumed quantity of product i (per capita), $q_{it} > 0$ is the quality of that product, $\theta > 1$ is the elasticity of substitution between products, $\mathcal{I}_t$ is the set of available products, and its (endogenous) cardinality is denoted by $M_t \equiv |\mathcal{I}_t|$.[3]

Each of those products is produced by a single firm using production labor according to the linear technology:

$$y_{it} = l_{it}^{\mathrm{P}}, \tag{4}$$

where y_{it} denotes the quantity of product i supplied at time t by the firm and l_{it}^{P} denotes the quantity of labor used in production.

Over time, a firm can incrementally improve the quality of its product by directing

[3]By available products, we mean products that have been invented and whose quality has not yet fallen below the exit threshold defined below.

labor towards development. More precisely, a product's quality evolves according to the following controlled process:

$$\dot{q}_{it} = \gamma_{it} \cdot q_{it}, \tag{5}$$

where γ_{it} is the proportional drift of the product's quality. By devoting $l_{it}^{\mathtt{D}}$ units of labor to development the firm can achieve a quality drift of:

$$\gamma_{it} = \left[\frac{(1+\zeta) l_{it}^{\mathtt{D}}}{c_{\mathtt{D}} (q_{it}/Q_t)^{\theta-1}} \right]^{\frac{1}{1+\zeta}}, \tag{6}$$

where $c_{\mathtt{D}} > 0$ determines the scale of the corresponding development cost function, $\zeta > 0$ measures its elasticity, and Q_t is an average quality index defined as:

$$Q_t \equiv \left(M_t^{-1} \int_{i \in \mathcal{I}_t} q_{it}^{\theta-1} \mathrm{d}i \right)^{\frac{1}{\theta-1}}.$$

Following Lashkari (2023), the development technology is such that there is a finite upper bound $\overline{\gamma} > 0$ on the drift of product quality to ensure that the optimal allocation does not involve any corner solutions as described in Trammell (2025).[4]

However, at Poisson rate $\epsilon > 0$, a firm may receive an idiosyncratic "obsolescence" shock after which it can no longer improve its product's quality. This is an extreme and stylized case of "ideas becoming harder to find" (Bloom et al., 2020) in which innovation efficiency literally falls to zero once the shock hits. After being hit by this shock, a product's quality eventually falls below a certain threshold $\underline{q}_t = \underline{q} \cdot Q_t$ where $\underline{q} \in (0,1)$, and the associated firm exogenously exits the market.

In every point in time, a unit measure of potential entrants attempt to discover products that are entirely new to society through research. Specifically, these entrants can direct $c_{\mathtt{R}}$ units of labor to research in order to invent a unit flow of these new products. Once a product is discovered, its initial quality equals the lower bound $\underline{q}_t$ of the product quality support.

Labor, which is inelastically supplied by the household, can be allocated to either production, development, or research, delivering the labor resource constraint:

$$L_t^{\mathtt{P}} + L_t^{\mathtt{D}} + L_t^{\mathtt{R}} \leq N_t, \tag{7}$$

[4] We thank Phillip Trammell and Chad Jones for pointing this out to us. We choose the value of $\overline{\gamma}$ such that this constraint does not bind in the decentralized equilibrium. This delivers a restriction on parameters such that $\overline{\gamma}$ must be greater than the corresponding expression in (11) below.

where these aggregate labor allocations are defined as:

$$L_t^{\mathrm{P}} \equiv \int_{i\in\mathcal{I}_t} l_{it}^{\mathrm{P}}\mathrm{d}i, \quad L_t^{\mathrm{D}} \equiv \int_{i\in\mathcal{I}_t} l_{it}^{\mathrm{D}}\mathrm{d}i, \quad \text{and} \quad L_t^{\mathrm{R}} \equiv c_{\mathrm{R}}R_t,$$

where R_t is the flow of new products discovered at time t. Finally, the resource constraint for each product is given by:

$$c_{it}N_t \le y_{it}, \ \ \forall i \in \mathcal{I}_t. \tag{8}$$

2.2 The decision problems

In this section, we define the decision problems of each economic agent. In terms of market structure, we assume that agents partake in perfect competition in all markets besides firms who engage in monopolistic competition.

The household's problem

Taking prices as given, the household's problem is to choose its consumption of each product to maximize lifetime utility:

$$\max_{\{a_t, c_t, \{c_{it}\}_{i\in\mathcal{I}_t}\}_{t\ge 0}} \int_0^\infty e^{-(\rho-n)t}\ln(c_t)\mathrm{d}t$$

subject to (3), a standard no-Ponzi game condition, and the flow budget constraint:

$$\dot{a}_t + \int_{i\in\mathcal{I}_t} p_{it}c_{it}\mathrm{d}i \le (r_t - n)a_t + w_t - T_t.$$

Here, p_{it} is the price of product i, w_t is the wage rate, T_t are lump-sum per-capita taxes, a_t is the value of corporate assets per capita, and r_t is the rate of return on those assets:

$$a_tN_t = \int_{i\in\mathcal{I}_t} V_{it}\mathrm{d}i, \quad \text{where} \quad \lim_{t\to\infty} e^{-\int_0^t (r_{t'}-n)\mathrm{d}t'}a_t = 0,$$

and where V_{it} denotes the value of product i at time t. This problem thus delivers the usual intertemporal Euler equation:

$$\dot{c}_t = (r_t - \rho)c_t,$$

and the per-capita demand functions:

$$c_{it} = (P_t/p_{it})^\theta q_{it}^{\theta-1}c_t. \tag{9}$$

In what follows, we will denote aggregate consumption by $C_t = c_t N_t$. Since aggregate consumption is chosen as the numéraire, the price index P_t is normalized to one for all t:

$$P_t \equiv \left(\int_{i \in \mathcal{I}_t} (p_{it}/q_{it})^{1-\theta} \mathrm{d}i \right)^{\frac{1}{1-\theta}} = 1.$$

The firm's problem

Firms engage in monopolistic competition in the product market but perfect competition in the labor market, taking the wage and the demand for their product as given. A firm thus chooses the price at which to sell its product and its labor demands to maximize the expected present discounted value of its profits.

From this point on, we abandon the i-index notation since a firm is entirely described by its product's quality q and its obsolescence status. We denote the latter by $\mathtt{S} \in \{\mathtt{O}, \mathtt{N}\}$ where $\mathtt{O}$ represents an "old" firm that has received the obsolescence shock, and $\mathtt{N}$ represents a "new" firm that has not. The new firm's value function satisfies a standard Hamilton-Jacobi-Bellman (HJB) equation:

$$\begin{aligned} r_t V_t^{\mathtt{N}}(q) = \max_{\mathbf{u}_t^{\mathtt{N}}(q) \geq \mathbf{0}} & \left\{ (1-\tau^{\mathtt{C}})[p_t(q) y_t(q) - w_t l_t^{\mathtt{P}}(q)] - (1-\tau^{\mathtt{D}}) w_t l_t^{\mathtt{D}}(q) + \gamma_t(q) q \partial_q V_t^{\mathtt{N}}(q) \right\} \\ & + \epsilon [V_t^{\mathtt{O}}(q) - V_t^{\mathtt{N}}(q)] + \dot{V}_t^{\mathtt{N}}(q) \end{aligned}$$

subject to the production technology (4), the development technology (6), the product resource constraint (8), and the demand function (9). Here, the vector $\mathbf{u}_t^{\mathtt{N}}(q) \equiv \{p_t(q), l_t^{\mathtt{P}}(q), l_t^{\mathtt{D}}(q)\}$ summarizes the control variables, $\tau^{\mathtt{C}} \geq 0$ is the corporate income tax rate, and $\tau^{\mathtt{D}} \geq 0$ is a subsidy on the firm's development expenditures. Similarly, the "old" firm's value function satisfies the HJB equation:[5]

$$r_t V_t^{\mathtt{O}}(q) = \max_{\mathbf{u}_t^{\mathtt{O}}(q) \geq \mathbf{0}} \left\{ (1-\tau^{\mathtt{C}})[p_t(q) y_t(q) - w_t l_t^{\mathtt{P}}(q)] \right\} + \dot{V}_t^{\mathtt{O}}(q)$$

subject to the production technology (4), the product resource constraint (8), and the demand function (9), and where $\mathbf{u}_t^{\mathtt{O}}(q) \equiv \{p_t(q), l_t^{\mathtt{P}}(q)\}$. The profit-maximization problem implies that a firm sets its price at a constant markup above marginal cost irrespective of its obsolescence status:

$$p_t(q) = \mu \cdot w_t, \ \ \forall q, t, \quad \text{where} \quad \mu \equiv \frac{\theta}{\theta - 1}.$$

[5]At the exit threshold, we have a boundary condition reflecting exit as detailed in Appendix A.1.

Hence, in status "S", a firm's flow profits can be expressed as:

$$\pi_t^{\mathtt{S}}(q) = (1-\tau^{\mathtt{C}})(q/Q_t)^{\theta-1}C_t/(\theta M_t) - \mathbb{1}_{\{\mathtt{S}=\mathtt{N}\}}(1-\tau^{\mathtt{D}})w_t l_t^{\mathtt{D}}(q),$$

where the optimal quality drift is:

$$\gamma_t(q) = \left[\frac{q\partial_q V_t^{\mathtt{N}}(q)}{(1-\tau^{\mathtt{D}})w_t c_{\mathtt{D}}(q/Q_t)^{\theta-1}}\right]^{1/\zeta}.$$

Substituting this into (6) and solving for $l_{it}^{\mathtt{D}}$ delivers the optimal firm-level allocation of labor to development.

The entrant's problem

Entrants engage in perfect competition on the labor market and, thus, choose research labor $L_t^{\mathtt{R}}$ to maximize future expected profits while taking the wage rate as given:

$$V_t^{\mathtt{E}} = \max_{L_t^{\mathtt{R}}} \left\{ V_t^{\mathtt{N}}(\underline{q}_t)L_t^{\mathtt{R}}/c_{\mathtt{R}} - (1-\tau^{\mathtt{R}})w_t L_t^{\mathtt{R}} \right\},$$

where $\tau^{\mathtt{R}} \geq 0$ is a research subsidy. The first-order condition of the entrant's problem delivers what will be referred to as the free-entry condition:

$$V_t^{\mathtt{N}}(\underline{q}_t) = (1-\tau^{\mathtt{R}})w_t c_{\mathtt{R}}.$$

2.3 The market equilibrium allocation

Having defined the decision problems of each economic agent, we can now define the concept of a market equilibrium allocation, and lay out the equations that determine the long-run equilibrium growth rate of the aggregate economy.

Definition 1. *Given the initial conditions $\{N_0, Q_0, \{m_0^{\mathtt{N}}(q), m_0^{\mathtt{O}}(q)\}_{q=\underline{q}_0}^{\infty}\}$, where $m_0^{\mathtt{S}}(q)$ is the initial measure of type-*S *firms with product quality q, a market allocation consists of time paths for quantities and prices such that the following conditions hold:*

1. $\{\{c_t(q)\}_{q=\underline{q}_t}^{\infty}\}_{t\geq 0}$ *solve the household's problem.*
2. $\{\{p_t(q), l_t^{\mathtt{P}}(q), l_t^{\mathtt{D}}(q)\}_{q=\underline{q}_t}^{\infty}\}_{t\geq 0}$ *solve the firm's problem.*
3. $\{L_t^{\mathtt{R}}\}_{t\geq 0}$ *solve the entrant's problem.*

4. $\{\{p_t(q)\}_{q=\underline{q}_t}^{\infty}\}_{t\geq 0}$ *clear the product markets.*

5. $\{w_t\}_{t\geq 0}$ *clear the labor market.*

6. $\{r_t\}_{t\geq 0}$ *clear the asset market.*

7. *The government's budget is balanced:* $T_t = \tau^{\mathtt{R}} w_t L_t^{\mathtt{R}} + \tau^{\mathtt{D}} w_t L_t^{\mathtt{D}} - \tau^{\mathtt{C}} C_t/\theta$.

The market allocation in this model admits a remarkably simple aggregation such that aggregate consumption is given by:[6]

$$C_t = M_t^{\frac{1}{\theta-1}} Q_t L_t^{\mathtt{P}}.$$

That is, aggregate consumption is increasing in the measure of products (owing to a taste for variety), the average quality across those products, and the labor input used in production. In Appendix A.1 we show that on a balanced growth path (BGP), the measure of products grows at the same rate as the population, the average quality index grows at a constant rate, and, thus, the growth rate of consumption per capita is constant and given by:

$$g = \frac{n}{\theta - 1} + g^Q. \tag{10}$$

In this expression, the growth rate of the average quality index g^Q is:

$$g^Q = \gamma - \frac{(n+\epsilon)(1-\underline{q}^{\theta-1})}{\theta - 1} \quad \text{where} \quad \gamma = \left[\frac{(\theta-1)(1-\tau^{\mathtt{R}})c_{\mathtt{R}}}{(1-\tau^{\mathtt{D}})c_{\mathtt{D}}\underline{q}^{\theta-1}}\right]^{1/\zeta} \tag{11}$$

which only depends on exogenous parameters and policy instruments, making clear that long-run growth can be sustained even under a constant population ($n = 0$).[7] When $n > 0$, an equivalent but more intuitive expression for g^Q is:

$$g^Q = \gamma \cdot \frac{n+d}{n+\epsilon} - \frac{n(1-\underline{q}^{\theta-1})}{\theta - 1} \tag{12}$$

where $d = \epsilon g^Q/\gamma$ is the endogenous exit rate. The first term is the product of the common product quality drift γ among firms that haven't yet received the obsolescence shock and the fraction of such firms in the economy $\frac{n+d}{n+\epsilon}$.[8] The second term reflects the

[6]See Appendix A.1.

[7]Without population growth, the economy still features *gross* entry and exit: obsolete firms hitting the exit threshold are replaced one-for-one by new entrants with renewed scope for development.

[8]Note that it is an equilibrium outcome that the product quality drift γ is constant over time and common across all firms in the `N` state.

negative growth contribution of net product entry occurring at rate n. Indeed, entering firms draw their initial product quality at the lower bound $\underline{q}_t$ of the quality support, which drags down the average quality index.

Let us now consider the expression for γ, which is where market incentives and R&D policy have their bite. First, a subsidy (τ^{D}) to development expenditures directly increases γ by reducing the cost of improving a product's quality. In contrast, a research subsidy (τ^{R}) lowers γ as it encourages entry, intensifying competition, and eroding expected profits per firm. This, in turn, dampens firms' incentives to invest in development. Finally, γ is neutral to corporate income taxes (τ^{C}) in the long run. While higher corporate taxes reduce post-tax profits, they also make entry less attractive. Through free entry, the number of firms in the market endogenously adjusts until expected profits per firm are restored, as discussed in Peretto (1998).

The fraction $\frac{n+d}{n+\epsilon}$ of non-obsolete firms is itself endogenous. It is decreasing in the arrival rate of obsolescence shocks ϵ and increasing in the endogenous exit rate $d = \epsilon g^{Q}/\gamma$ of firms. This exit rate is increasing in g^{Q}: faster growth of the average quality index intensifies the race against the exit threshold, so obsolete firms hit that threshold sooner. By contrast, it is decreasing in γ: faster product-quality growth raises the likelihood that firms have upgraded enough before receiving the obsolescence shock to move away from the threshold.

Finally, these R&D and obsolescence forces also pin down the cross-sectional distribution of product quality. Indeed, the stationary firm size distribution is Pareto with shape parameter $\lambda = \frac{1}{1-\underline{q}^{\theta-1}} > 1$. This distribution displays a fatter right tail when the initial quality of entrants $\underline{q}$ is lower (close to zero). Conversely, when entrants are nearly as good as incumbents (i.e., $\underline{q}$ close to one), the distribution becomes thinner-tailed. This is because a lower initial quality of entrants implies that incumbents have more time to improve their products before facing comparable competition from new entrants, leading to a wider dispersion in product quality and, consequently, firm size.

3 Discussion

In this section, we discuss the crucial assumptions that underlie our model and possible extensions to our framework.

3.1 Key assumptions

The predictions of our model hinge on two important assumptions. First, we assume that entrants draw an initial level of product quality that is proportional to the contemporaneous average productivity index Q_t. This proportionality is what sustains quality-led growth on a BGP. Without it (i.e., if entrants started producing at a fixed level of product quality), average quality growth would eventually stall. Although individual incumbent firms might temporarily improve their products' quality, these gains would be lost when they exit. New entrants would consistently reset the quality level to the same starting point, preventing any cumulative progress. The model would then revert to a semi-endogenous growth framework where per-capita consumption growth would be fueled exclusively by the introduction of new products, through a taste for variety.

This first assumption makes our model consistent with evidence from Garcia-Macia et al. (2019) who find that about three quarters of productivity growth in the U.S. between 2003 and 2013 occurred through incumbents' improvements to their existing products rather than the introduction of entirely new products. Therefore, without this quality-scaling at entry, our model would attribute all of long-run growth to net entry and none of it to incremental improvements, contrary to their findings. Jensen, McGuckin and Stiroh (2001) also provide evidence that more recent cohorts of entrants among U.S. manufacturing plants have higher productivity than earlier cohorts, which is consistent with the assumption that entrants are "building on the shoudlers of giants".

The second key assumption is that entry costs are constant in units of labor and therefore scale in the wage rate which, in equilibrium, grows at the same rate as firm-level profits. This keeps the free-entry margin "balanced": if entry costs grew faster than firm-level profits, entry would be choked off, the firm count would lag population growth, and development labor per firm would rise over time, delivering explosive growth with strong scale effects. This outcome would be at odds with empirical evidence and with our objective of neutralizing strong scale effects. As in second-generation endogenous growth models, by tying entry costs to the wage, the number of firms can track population, and development effort per firm remains constant.

Two sets of facts support this assumption. Direct evidence in Klenow and Li (2024) indicates that revenue per firm in the U.S. increases with the level of productivity (both over time and across states). If higher revenues are associated with higher profits, entry costs must also rise with productivity to satisfy the free-entry condition. Therefore, they conclude that entry costs rise with growth, as in our model. Complementarily, Laincz and Peretto (2006) document that, in the U.S., the number of firms scales with population and the number of R&D workers per firm is roughly stable over time. In Appendix Figures B.4 and B.5, we show that the distribution of employment across establishments is also roughly stationary using data from the U.S. Business Dynamics Statistics. Together, these observations justify our treatment of entry costs and underscore why this assumption is central to ruling out strong scale effects in the model.

3.2 Potential extensions

Our baseline model is deliberately parsimonious. Here, we outline several directions in which it could be extended. First, we make several simplifying assumptions regarding the evolution of product quality. Equation (5) specifies a deterministic process for product quality but one could instead allow for stochastic dynamics (e.g., Brownian shocks or Poisson jumps), which would then deliver additional cross-sectional dispersion. Equation (6) scales development costs by $(q_{it}/Q_t)^{\theta-1}$ to ensure tractability and consistency with Gibrat's law among firms in the "N" state, but it is not a necessary condition for our main mechanism to operate. We also model technological obsolescence as a Poisson shock that shuts down the possibility for further development, but one could instead assume the development efficiency in a line to decline more gradually with product quality. Moreover, our CES demand specification implies constant markups, but one could extend our model to deliver firm-level markup heterogeneity through a non-CES demand system or Bertrand competition as in Peters (2020).

Second, our stylized assumptions regarding entry and exit could be relaxed. We impose an exogenous exit threshold $\underline{q}_t = \underline{q} \cdot Q_t$, but we could instead replace it by an endogenous exit rule in the presence of overhead costs. We assume that entrants draw their initial quality at that exit threshold $\underline{q}_t$, but one could entertain alternative learning processes for new entrants (e.g., Lashkari 2023; Yao 2024). These modifications would affect the contributions of entry and exit to quality growth in equation (11), but would not otherwise affect our main model mechanism. Indeed, allowing entry and exit to contribute positively to average quality growth (either directly or indirectly through technology spillovers) would accelerate the displacement of laggards such that the model would feature more creative destruction.

Third, we make simplifying assumptions about the boundary of the firm and the allocation of R&D activities. While we treat firms as single-product entities for tractability, our central mechanism could be embedded in a multi-product environment as in Klette and Kortum (2004) and Peters and Walsh (2021) and allow incumbents to engage in research as well.[9] Such an extension would have notable policy implications. In our model, research and development can be targeted differentially because they are conducted by distinct agents: an entry subsidy targets research separately from a de-

[9] See also Akcigit and Kerr (2018), who distinguish between incremental and radical innovation.

velopment subsidy. In contrast, if both research and development were performed by incumbents, it may prove difficult to differentially subsidize these two activities. In that case, setting a unique subsidy rate for R&D ($\tau^R = \tau^D$) would leave long-run growth unaffected, as can be seen from (11).[10]

4 A simple calibration

In this section, we present a straightforward calibration of our model's parameters to illustrate its mechanisms. Although this exercise is not meant as a rigorous empirical quantification for direct comparison with economic data, it provides insight into the potential magnitude of different forces.

We set the pure rate of time preference, ρ, to 0.04 and assume an annual population growth rate, n, of 1%. Consistent with Garcia-Macia et al. (2019), the elasticity of substitution across products, θ, is set to 4. We fix ζ, which captures the degree of decreasing returns to development labor, at 1, which aligns with the preferred value in Acemoglu, Akcigit, Alp, Bloom and Kerr (2018).

Four additional parameters require calibration: the development cost parameter c_{D}, the research cost parameter c_{R}, the obsolescence shock arrival rate ϵ, and the initial relative quality of new products $\underline{q}$. Since these parameters are less conventional in the literature or not directly observable, we calibrate them by jointly matching a growth rate of per capita consumption of 2% per year, as well as the following three establishment-level moments, which we calculate from the U.S. Business Dynamics Statistics (BDS) between 2015 and 2019:[11]

1. The (Pareto) tail index of the establishment size distribution is 2.95.[12]
2. The establishment-level exit rate is 8.7%.
3. The average establishment has 18.1 employees.

The calibrated parameter values are reported in Table 1 and all four of the above empirical moments are matched exactly.

[10] In this case, our model's policy predictions essentially coincide with those of semi-endogenous growth frameworks.

[11] We use establishment- rather than firm-level data since every firm produces a single product in our model. We view a firm as a set of products and proxy products by establishments, following Garcia-Macia et al. (2019).

[12] The distribution of employment across U.S. establishments is better approximated by a lognormal distribution. Therefore, we fit a lognormal distribution and calculate the local Pareto tail index above 250 employees.

Table 1: Calibration

Parameter	Value	Source
ρ	0.04	Standard
n	0.01	Population growth
θ	4	Garcia-Macia et al. (2019)
ζ	1	Acemoglu et al. (2018)
$c_{\mathtt{D}}$	exp(8.55)	Per capita consumption growth of 2%
$c_{\mathtt{R}}$	exp(3.90)	Average establishment employment of 18.1 (U.S. BDS)
$\underline{q}$	0.87	Pareto tail index of 2.95 (U.S. BDS)
ϵ	0.23	Establishment exit rate of 8.7% (U.S. BDS)

While the parameters $\{c_{\mathtt{D}}, c_{\mathtt{R}}, \epsilon, \underline{q}\}$ are jointly calibrated to match the aforementioned moments, we provide some intuition for their identification. The development cost parameter $c_{\mathtt{D}}$ is primarily identified by the growth rate of per capita consumption, as it governs the pace at which firms improve their product's quality and, in turn, economic growth. The research cost parameter $c_{\mathtt{R}}$ is pinned down by the average establishment size, since it influences the cost of entry and thereby the equilibrium number of firms. The obsolescence shock arrival rate, ϵ, is identified from the establishment-level exit rate because it governs how quickly products become obsolete and ultimately exit the market. Finally, as discussed in Section 2.3, given θ, the initial relative quality of new products, $\underline{q}$, directly determines the Pareto tail index of the establishment size distribution.

Growth accounting. Under this simple calibration, the long-run growth rate of per-capita consumption can be decomposed into a semi-endogenous component from variety expansion and an endogenous component from quality improvements. The latter can be split further into (1) the contribution of quality improvements by incumbents and (2) a dilution term from net entry, reflecting that entrants start below the contemporaneous average quality. Table 2 reports this decomposition.

Table 2: Sources of economic growth

Source	Contribution (p.p.)	Share of total
Variety growth (semi-endogenous)	0.33%	16.7%
Quality growth (endogenous)	1.67%	83.3%
Development (incumbents)	1.78%	89.0%
Research (net entry)	-0.11%	-5.7%
Total	2.00%	100%

Variety expansion contributes 0.33 percentage points, or 16.7% of total growth. The remaining 1.67 percentage points (83.3%) come from endogenous quality growth. Within this component, incumbent development contributes 1.78 percentage points, while net entry subtracts 0.11 percentage points. While turnover sustains long-run growth by replenishing the pool of product lines to improve, aggregate quality growth comes in this accounting sense mainly from within-product improvements by incumbents.

This important role played by incumbents is consistent with Garcia-Macia et al. (2019) and Peters and Walsh (2021), who attribute the bulk of U.S. productivity growth to incumbent innovation on existing products, with a smaller role for net variety creation. Our estimate of the incumbent share (89%) is somewhat higher than these empirical findings (typically 60–75%). This difference is largely mechanical: to preserve analytical tractability, we assume entrants enter at the exit threshold, which is significantly below the average quality, creating a negative dilution effect. If we instead assumed entrants began with higher initial quality—reflecting, for instance, a more realistic learning process as in Yao (2024) or Lashkari (2023)—the contribution of net entry would rise, and that of incumbent innovation would correspondingly decline.

Because we have free entry in our model, entry and exit rates are increasing in the population growth, echoing Peters and Walsh (2021)'s finding that the demographic slowdown accounts for much of the decline in business dynamism. Hence, as population growth declines, our model is consistent with the secular decline in U.S. business dynamism documented by Decker, Haltiwanger, Jarmin and Miranda (2014). Crucially, Garcia-Macia et al. (2019) also show that the recent U.S. TFP growth slowdown results almost entirely from weaker creative destruction and slower variety creation—not from reduced incumbent innovation on existing products. Since product development by incumbents in our model is insulated from demographics while variety expansion is not, slowing population growth affects the endogenous component of growth less while it dampens the semi-endogenous component proportionally.

4.1 Policy experiments

Using this calibration, we can conduct simple policy experiments to illustrate the effects of different R&D and tax policies. We consider two types of policies: a 10% subsidy to development expenditures ($\tau^{\mathtt{D}}$) and a 10% subsidy to research expenditures ($\tau^{\mathtt{R}}$). In Appendix B, we also consider a 10% cut in the corporate income tax rate ($\tau^{\mathtt{C}}$) which leaves long-run growth unaffected.

Development subsidy. Figure 1 shows the effect of a 10% subsidy to development expenditures on the growth rate of consumption per capita (in percentage points). This policy has a powerful positive effect on long-run growth, increasing it by almost 50 basis points.[13] This is because it directly encourages firms to invest in improving their products, which raises the growth rate of consumption per capita through equation (10). In the short run, the policy generates a temporary growth dip of about 4 percentage points as labor is reallocated away from production and research, reducing consumption and the variety of products available to consumers. The figure shows a brief positive spike at the moment of implementation, reflecting the instantaneous jump in the common product quality growth rate.

Figure 1: 10% development subsidy

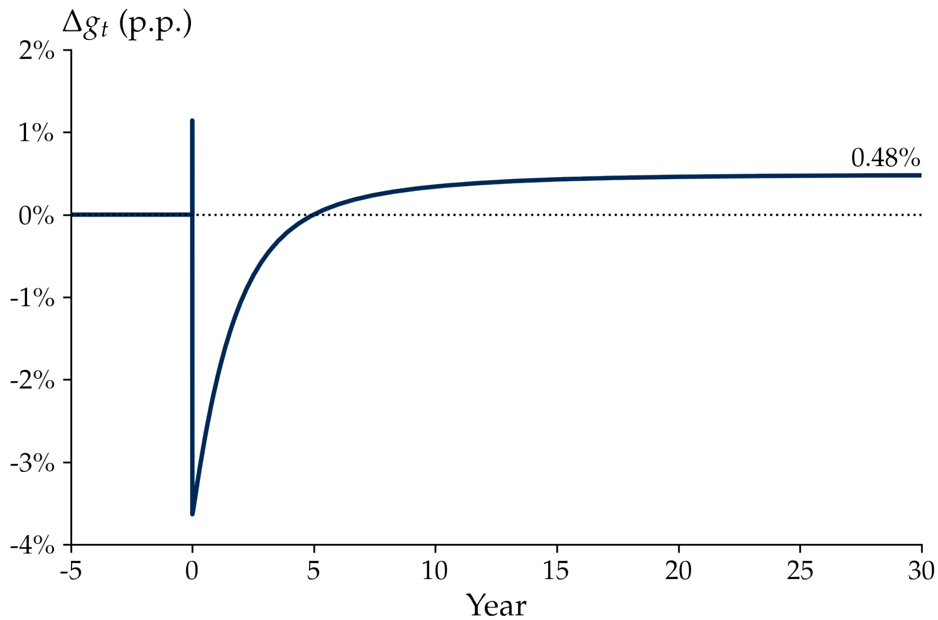

Research subsidy. Figure 2 reports the growth effect of a 10% research subsidy. This policy reduces long-run consumption per capita growth by roughly 43 basis points, as it promotes entry and competition, which erodes profits and discourages quality-improving development. The figure shows a brief negative spike of about 6 percentage points at the moment of implementation, reflecting the instantaneous slowdown in the common product quality growth rate. This is followed by a positive bounce as product variety expands. It is worth noting that our model abstracts from the possibility that research could deliver positive technological spillovers, in which case subsidies may enhance long-run growth.

[13] While our model is intentionally parsimonious and the magnitude of this effect should be interpreted with caution, it is broadly consistent with growth effect estimates reported elsewhere in the literature, such as Acemoglu et al. (2018) and Aghion, Bergeaud, Boppart, Klenow and Li (2025).

Figure 2: 10% research subsidy

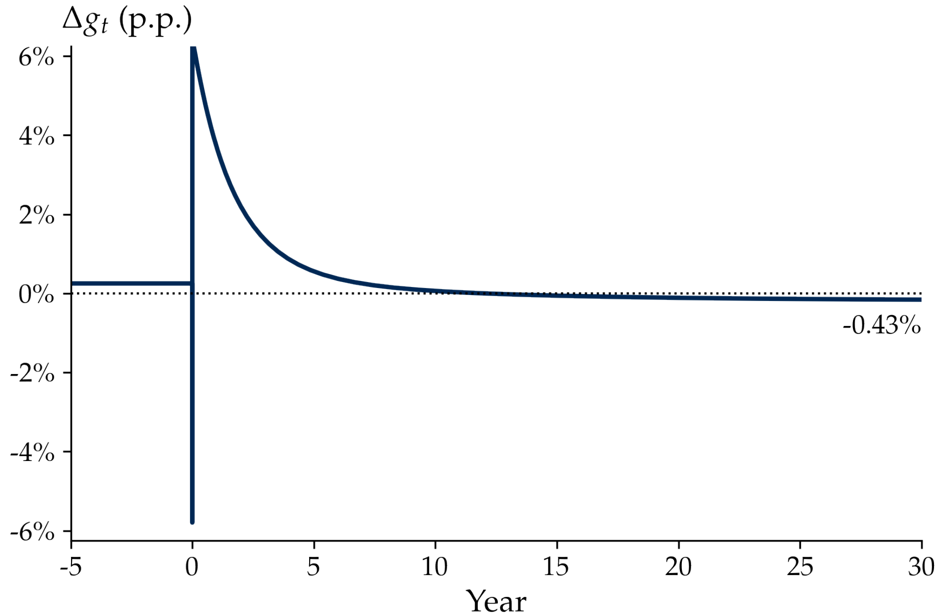

These experiments highlight the different transitional and long-run consequences of R&D policies. They should be interpreted as stylized illustrations, and it is important to note that, in our model, the market equilibrium allocation is constrained-efficient, as shown in Appendix A.2. Thus, none of these interventions would improve welfare once we abstract from technological spillovers. As such, the welfare relevance of these policies ultimately depends on forces that lie outside the model's scope.

5 Conclusion

In this paper, we propose a new framework for endogenous economic growth that captures two key empirical facts: the absence of strong scale effects and the evidence that ideas are becoming harder to find at the micro level. Our model distinguishes between research (the creation of new products) and development (the improvement of existing products). Diminishing returns to development cause progress upon any given product to eventually stall. However, the continuous entry of new products via research "resets the innovation clock," creating a perpetual stream of new development opportunities that sustains aggregate growth.

This turnover-driven mechanism allows the model to generate a constant, endogenous rate of long-run growth while remaining consistent with the evidence of diminishing returns to innovation at both the macro and micro levels. Our calibration, disciplined by U.S. establishment-level data, demonstrates that policies directly targeting incumbent innovation, such as development subsidies, can have a powerful positive effect on long-run growth, whereas policies encouraging entry can backfire by intensifying product-market competition and weakening development incentives.

Beyond the aggregate moments matched in our simple calibration exercise, our model delivers sharp, testable predictions about product and firm dynamics—in particular, how sales and employment growth evolve over the firm's lifecycle and how the distributions of firm ages and sizes respond to different policies. A productive avenue for future research would be to confront these richer predictions with granular microdata on products and establishments in order to flesh out the bones of our model and uncover how the "innovation clock" resets empirically.

Finally, the implications of our framework are particularly important in light of the projected slowdown in global population growth. A standard semi-endogenous growth model, where long-run growth is crucially determined by population growth, would predict "the end of economic growth" as the latter falls to zero (Jones, 2022). Our model, in contrast, offers a more optimistic outlook. Because long-run growth is partially driven by a policy-sensitive engine of quality upgrading insulated from population dynamics, our framework shows that sustained growth is possible even with a constant population. By providing a theory of endogenous growth that is consistent with the key empirical regularities of the U.S. economy, our model suggests that the future of economic growth need not be bleak.

References

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr, "Innovation, Reallocation, and Growth," *American Economic Review*, November 2018, *108* (11), 3450–91.

Aghion, Philippe and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica*, 1992, *60* (2), 323–351.

_ and _ , "Research and Development in the Growth Process," *Journal of Economic Growth*, 1996, *1*, 49–73.

_ , Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li, "Good Rents versus Bad Rents: R&D Misallocation and Growth," Working Paper 34190, National Bureau of Economic Research August 2025.

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

_ , John Grigsby, Tom Nicholas, and Stefanie Stantcheva, "Taxation and Innovation in the Twentieth Century," *The Quarterly Journal of Economics*, 06 2021, *137* (1), 329–385.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?," *American Economic Review*, April 2020, *110* (4), 1104–44.

Dechezleprêtre, Antoine, Elias Einiö, Ralf Martin, Kieu-Trang Nguyen, and John Van Reenen, "Do tax incentives increase firm innovation? an rd design for r&d, patents, and spillovers," *American Economic Journal: Economic Policy*, 2023, *15* (4), 486–521.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, "The Role of Entrepreneurship in US Job Creation and Economic Dynamism," *Journal of Economic Perspectives*, September 2014, *28* (3), 3–24.

Dhingra, Swati and John Morrow, "Monopolistic Competition and Optimum Product Diversity under Firm Heterogeneity," *Journal of Political Economy*, 2019, *127* (1), 196–232.

Dinopoulos, Elias and Peter Thompson, "Schumpeterian Growth Without Scale Effects," *Journal of Economic Growth*, 1998, *3* (4), 313–335.

Dixit, Avinash K. and Joseph E. Stiglitz, "Monopolistic Competition and Optimum Product Diversity," *The American Economic Review*, 1977, *67* (3), 297–308.

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

Gordon, Robert, *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*, Princeton University Press, 2017.

Howitt, Peter, "Steady Endogenous Growth with Population and R & D Inputs Growing," *Journal of Political Economy*, 1999, *107* (4), 715–730.

Jensen, J. Bradford, Robert H. McGuckin, and Kevin J. Stiroh, "The Impact of Vintage and Survival on Productivity: Evidence from Cohorts of U.S. Manufacturing Plants," *The Review of Economics and Statistics*, 05 2001, *83* (2), 323–332.

Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, 1995, *103* (4), 759–784.

_ , "The End of Economic Growth? Unintended Consequences of a Declining Population," *American Economic Review*, November 2022, *112* (11), 3489–3527.

Klenow, Peter J. and Huiyu Li, "Entry Costs Rise with Growth," Working Paper 32974, National Bureau of Economic Research September 2024.

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

Laincz, Christopher and Pietro Peretto, "Scale Effects in Endogenous Growth Theory: an Error of Aggregation not Specification," *Journal of Economic Growth*, September 2006, *11* (3), 263–288.

Lashkari, Danial, "Innovation Policy in Theories of Knowledge Diffusion and Selection," Working Paper 2023.

Mokyr, Joel, "Punctuated equilibria and technological progress," *The American Economic Review*, 1990, *80* (2), 350–354.

_ , *The lever of riches: Technological creativity and economic progress*, Oxford University Press, 1992.

_ , "Secular stagnation? Not in your life," in Coen Teulings and Richard Baldwin, eds., *Secular Stagnation: Facts, Causes and Cures*, Paris & London: CEPR Press, 2014.

Peretto, Pietro F., "Technological Change and Population Growth," *Journal of Economic Growth*, 1998, *3* (4), 283–311.

Peters, Michael, "Heterogeneous Markups, Growth, and Endogenous Misallocation," *Econometrica*, 2020, *88* (5), 2037–2073.

_ **and Conor Walsh**, "Population Growth and Firm Dynamics," Working Paper 29424, National Bureau of Economic Research October 2021.

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

Segerstrom, Paul S., "Endogenous Growth without Scale Effects," *The American Economic Review*, 1998, *88* (5), 1290–1310.

Trammell, Philip, "Endogenous Growth and Excess Variety," Working Paper, Stanford University 2025.

U.S. Census Bureau, "Business Dynamics Statistics," 2023. Accessed: 2026.

Yao, Jinglun, "Knowledge is (Market) Power," Working Paper 2024.

Young, Alwyn, "Growth without Scale Effects," *Journal of Political Economy*, 1998, *106* (1), 41–63.

Appendix

A Theoretical appendix

This section of the appendix provides derivations and proofs for the results presented in the paper.

A.1 The market equilibrium allocation

Aggregate consumption. To derive an expression for aggregate consumption, we start from the definition of aggregate production labor:

$$L_t^{\mathtt{P}} = \int_{i\in\mathcal{I}_t} l_{it}^{\mathtt{P}} \mathrm{d}i.$$

Substituting in the firm's production function and the product resource constraints from (4) and (8), we obtain:

$$L_t^{\mathtt{P}} = N_t \int_{i\in\mathcal{I}_t} c_{it} \mathrm{d}i.$$

Substituting in the household's demand function from (9), we find:

$$L_t^{\mathtt{P}} = C_t \int_{i\in\mathcal{I}_t} q_{it}^{\theta-1} p_{it}^{-\theta} \mathrm{d}i.$$

Using the monopoly pricing condition ($p_{it} = \mu w_t$) and the definition of the aggregate quality index Q_t, we obtain:

$$L_t^{\mathtt{P}} = C_t (\mu w_t)^{-\theta} Q_t^{\theta-1} M_t. \tag{A.1}$$

Finally, the numéraire condition implies:

$$P_t \equiv \left(\int_{i\in\mathcal{I}_t} (p_{it}/q_{it})^{1-\theta} \mathrm{d}i \right)^{\frac{1}{1-\theta}} = 1 \quad \Leftrightarrow \quad \mu w_t = \left(\int_{i\in\mathcal{I}_t} q_{it}^{\theta-1} \mathrm{d}i \right)^{\frac{1}{\theta-1}} = M_t^{\frac{1}{\theta-1}} Q_t.$$

Substituting this result in (A.1), we obtain:

$$C_t = M_t^{\frac{1}{\theta-1}} Q_t L_t^{\mathtt{P}}.$$

Hamilton-Jacobi-Bellman equations. The old firm's value function satisfies the HJB equation:

$$r_t V_t^{\mathtt{O}}(q) = (1-\tau^{\mathtt{C}})(q/Q_t)^{\theta-1} c_t N_t / (\theta M_t) + \dot{V}_t^{\mathtt{O}}(q). \tag{A.2}$$

Defining $x_t \equiv \ln(q_t/Q_t)$, we can rewrite:

$$r_t V_t^{\mathtt{O}}(x) = (1-\tau^{\mathtt{C}})\exp[(\theta-1)x]c_t N_t/(\theta M_t) - g_t^Q \partial_x V_t^{\mathtt{O}}(x) + \dot{V}_t^{\mathtt{O}}(x).$$

We impose value matching at the exit threshold $\underline{q}_t = \underline{q} Q_t$. When a product line reaches $\underline{q}_t$, it is exogenously liquidated (or resold) and the proceeds are rebated lump-sum to households. Because all cash flows scale with relative quality, the competitive liquidation value inherits the same homotheticity. Therefore, under the conjecture that:

$$V_t^{\mathtt{O}}(x) = V_t^{\mathtt{O}} \exp[(\theta-1)x]$$

value matching implies the boundary condition:

$$V_t^{\mathtt{O}}(\underline{q}_t) = V_t^{\mathtt{O}} \underline{q}^{\theta-1}.$$

Substituting this guess into equation (A.2), we obtain the following ordinary differential equation (ODE):

$$\dot{V}_t^{\mathtt{O}} = [r_t + (\theta-1)g_t^Q]V_t^{\mathtt{O}} - (1-\tau^{\mathtt{C}})c_t N_t/(\theta M_t), \tag{A.3}$$

which verifies our guess. The new firm's value function satisfies the HJB equation:

$$\begin{aligned}(r_t+\epsilon)V_t^{\mathtt{N}}(q) = \max_{l_t^{\mathtt{D}}(q)} \{&(1-\tau^{\mathtt{C}})(q/Q_t)^{\theta-1}c_t N_t/(\theta M_t) - (1-\tau^{\mathtt{D}})w_t l_t^{\mathtt{D}}(q) \\ &+ \gamma_t(q) q \partial_q V_t^{\mathtt{N}}(q)\} + \epsilon V_t^{\mathtt{O}}(q) + \dot{V}_t^{\mathtt{N}}(q).\end{aligned}$$

Using the change of variable defined above, we can rewrite:

$$\begin{aligned}(r_t+\epsilon)V_t^{\mathtt{N}}(x) = \max_{l_t^{\mathtt{D}}(x)} \{&(1-\tau^{\mathtt{C}})\exp[(\theta-1)x]c_t N_t/(\theta M_t) - (1-\tau^{\mathtt{D}})w_t l_t^{\mathtt{D}}(x) \\ &+ [\gamma_t(x) - g_t^Q]\partial_x V_t^{\mathtt{N}}(x)\} + \epsilon V_t^{\mathtt{O}}(x) + \dot{V}_t^{\mathtt{N}}(x).\end{aligned} \tag{A.4}$$

Let us guess that this value function takes the following form:

$$V_t^{\mathtt{N}}(x) = V_t^{\mathtt{N}} \exp[(\theta-1)x]. \tag{A.5}$$

Substituting this guess into the optimal product quality drift, we obtain:

$$\gamma_t = \left[\frac{(\theta-1)V_t^{\mathtt{N}}}{(1-\tau^{\mathtt{D}})w_t c_{\mathtt{D}}}\right]^{1/\zeta}$$

which is independent of a product's quality. Substituting this result in (A.4), we obtain the following ODE, which verifies our guess:

$$\dot{V}_t^{\mathtt{N}} = [r_t + \epsilon - (\theta - 1)(\gamma_t - g_t^Q)]V_t^{\mathtt{N}} - \epsilon V_t^{\mathtt{O}} - \frac{(1-\tau^{\mathtt{C}})c_t N_t}{\theta M_t} + \frac{(1-\tau^{\mathtt{D}})w_t c_{\mathtt{D}} \gamma_t^{1+\zeta}}{1+\zeta}. \quad \text{(A.6)}$$

Fokker-Planck equations. The Fokker-Planck (FP) equations describing the evolution of the density of log relative quality among new and old firms are given by:

$$\dot{m}_t^{\mathtt{N}}(x) = -(\gamma_t - g_t^Q)\partial_x m_t^{\mathtt{N}}(x) - \epsilon m_t^{\mathtt{N}}(x),$$
$$\dot{m}_t^{\mathtt{O}}(x) = g_t^Q \partial_x m_t^{\mathtt{O}}(x) + \epsilon m_t^{\mathtt{N}}(x),$$

where $\underline{x} \equiv \ln(\underline{q}_t / Q_t)$, and we have the following boundary condition:

$$(\gamma_t - g_t^Q) \lim_{x \to \underline{x}} m_t^{\mathtt{N}}(x) = R_t.$$

Let us denote the aggregate measures of new and old products by:

$$M_t^{\mathtt{N}} \equiv \int_{\underline{x}}^{\infty} m_t^{\mathtt{N}}(x)\mathrm{d}x \quad \text{and} \quad M_t^{\mathtt{O}} \equiv \int_{\underline{x}}^{\infty} m_t^{\mathtt{O}}(x)\mathrm{d}x.$$

The laws of motion for these measures are given by:

$$\dot{M}_t^{\mathtt{N}} = R_t - \epsilon M_t^{\mathtt{N}} \quad \text{and} \quad \dot{M}_t^{\mathtt{O}} = \epsilon M_t^{\mathtt{N}} - d_t M_t$$

where d_t denotes the exit rate at the lower bound of the quality support:

$$d_t \equiv g_t^Q \lim_{x \to \underline{x}} m_t^{\mathtt{O}}(x) / M_t.$$

Hence, as we have $M_t = M_t^{\mathtt{N}} + M_t^{\mathtt{O}}$, the total measure of products evolves according to:

$$\dot{M}_t = (e_t - d_t)M_t$$

where $e_t \equiv R_t / M_t$ denotes the entry rate.

Equilibrium conditions. Using (A.5), the free-entry condition can be rewritten as follows:

$$V_t^{\mathtt{N}} \underline{q}^{\theta-1} = (1-\tau^{\mathtt{R}}) w_t c_{\mathtt{R}}. \quad \text{(A.7)}$$

The labor market clearing condition can be rewritten as:

$$N_t = \frac{c_t N_t}{\mu w_t} + \frac{c_{\mathtt{D}} \gamma_t^{1+\zeta}}{1+\zeta} \cdot \int_{\underline{x}}^{\infty} \exp[(\theta-1)x] m_t^{\mathtt{N}}(x) \mathrm{d}x + c_{\mathtt{R}} e_t M_t.$$

The only endogenous variable for which a corresponding equation is missing is the growth rate of the average quality index. Using the change of variable $x_t \equiv \ln(q_t/Q_t)$, the expression for this index implies:

$$M_t^{-1} \int_{\underline{x}}^{\infty} \exp[(\theta-1)x] m_t(x) \mathrm{d}x = 1, \quad \text{where} \quad m_t(x) \equiv m_t^{\mathtt{N}}(x) + m_t^{\mathtt{O}}(x).$$

Normalizations. Let us define normalized variables:

$$\mathcal{V}_t^{\mathtt{O}} \equiv \frac{V_t^{\mathtt{O}}}{c_t}, \quad \mathcal{V}_t^{\mathtt{N}} \equiv \frac{V_t^{\mathtt{N}}}{c_t}, \quad \mathcal{M}_t \equiv \frac{M_t}{N_t}, \quad \mathcal{S}_t^{\mathtt{N}} \equiv \frac{M_t^{\mathtt{N}}}{M_t}, \quad \mathcal{S}_t^{\mathtt{O}} \equiv \frac{M_t^{\mathtt{O}}}{M_t},$$

as well as the normalized distributions $f_t^{\mathtt{N}}(x) = m_t^{\mathtt{N}}(x)/M_t^{\mathtt{N}}$ and $f_t^{\mathtt{O}}(x) = m_t^{\mathtt{O}}(x)/M_t^{\mathtt{O}}$. With these definitions, we use the free-entry condition in equation (A.7) and the household's Euler equation to rewrite equations (A.3) and (A.6) as:

$$\dot{\mathcal{V}}_t^{\mathtt{O}} = [\rho + (\theta-1)g_t^{\mathrm{Q}}]\mathcal{V}_t^{\mathtt{O}} - \frac{1-\tau^{\mathtt{C}}}{\theta \mathcal{M}_t}, \tag{A.8}$$

$$\dot{\mathcal{V}}_t^{\mathtt{N}} = [\rho + \epsilon - (\theta-1)(\gamma - g_t^{\mathrm{Q}})]\mathcal{V}_t^{\mathtt{N}} - \epsilon \mathcal{V}_t^{\mathtt{O}} - \frac{1-\tau^{\mathtt{C}}}{\theta \mathcal{M}_t} + \frac{(\theta-1)\gamma \mathcal{V}_t^{\mathtt{N}}}{1+\zeta} \tag{A.9}$$

where the product quality drift can be rewritten as follows using the free-entry condition:

$$\gamma = \left[\frac{(\theta-1)(1-\tau^{\mathtt{R}})c_{\mathtt{R}}}{(1-\tau^{\mathtt{D}})c_{\mathtt{D}} \underline{q}^{\theta-1}} \right]^{1/\zeta}.$$

Hence, the product quality drift is constant along any transition path. Similarly, the labor market clearing condition can be rewritten as:

$$1 = \frac{(1-\tau^{\mathtt{R}})c_{\mathtt{R}}}{\mu \mathcal{V}_t^{\mathtt{N}} \underline{q}^{\theta-1}} + \frac{c_{\mathtt{D}} \gamma^{1+\zeta}}{1+\zeta} \cdot \mathcal{M}_t \mathcal{S}_t^{\mathtt{N}} \int_{\underline{x}}^{\infty} \exp[(\theta-1)x] f_t^{\mathtt{N}}(x) \mathrm{d}x + c_{\mathtt{R}} e_t \mathcal{M}_t.$$

The Fokker-Planck (FP) equations describing the evolution of the distribution of log relative quality among new and old firms are given by:

$$\dot{f}_t^{\mathtt{N}}(x) = -(\gamma - g_t^Q)\partial_x f_t^{\mathtt{N}}(x) - (e_t/\mathcal{S}_t^{\mathtt{N}})f_t^{\mathtt{N}}(x),$$
$$\dot{f}_t^{\mathtt{O}}(x) = g_t^Q \partial_x f_t^{\mathtt{O}}(x) + (\epsilon \mathcal{S}_t^{\mathtt{N}}/\mathcal{S}_t^{\mathtt{O}})[f_t^{\mathtt{N}}(x) - f_t^{\mathtt{O}}(x)] + (d_t/\mathcal{S}_t^{\mathtt{O}})f_t^{\mathtt{O}}(x)$$

where $(\gamma - g_t^Q)\lim_{x\to\underline{x}} f_t^{\mathtt{N}}(x) = e_t/\mathcal{S}_t^{\mathtt{N}}$ and $d_t \equiv g_t^Q \mathcal{S}_t^{\mathtt{O}} \lim_{x\to\underline{x}} f_t^{\mathtt{O}}(x)$. The law of motion for the share of new and old products are given by:

$$\dot{\mathcal{S}}_t^{\mathtt{N}} = e_t - \mathcal{S}_t^{\mathtt{N}}(e_t + \epsilon - d_t) \quad \text{and} \quad \mathcal{S}_t^{\mathtt{O}} = 1 - \mathcal{S}_t^{\mathtt{N}},$$

and the law of motion for the measure of products per capita is given by:

$$\dot{\mathcal{M}}_t = (e_t - d_t - n)\mathcal{M}_t.$$

Finally, the expression for the average quality index implies:

$$\int_{\underline{x}}^{\infty} \exp[(\theta - 1)x][f_t^{\mathtt{N}}(x)\mathcal{S}_t^{\mathtt{N}} + f_t^{\mathtt{O}}(x)\mathcal{S}_t^{\mathtt{O}}]\mathrm{d}x = 1. \tag{A.10}$$

Balanced growth path. On a BGP, the growth rate of the average quality index and the drift of product quality for new firms are both constant. Moreover, the measure of new and old firms both grow at the same rate as the population. This implies that the entry and exit rates are constant, and the former is equal to $e = n + d$. The share of new and old products are also constant and equal to:

$$\mathcal{S}^{\mathtt{N}} = \frac{n+d}{n+\epsilon} \quad \text{and} \quad \mathcal{S}^{\mathtt{O}} = \frac{\epsilon - d}{n+\epsilon}. \tag{A.11}$$

The stationary FP equation for the distribution of log relative quality among new firms is given by:

$$-(\gamma - g^Q)\partial_x f^{\mathtt{N}}(x) - (n+\epsilon)f^{\mathtt{N}}(x) = -(n+\epsilon)\delta(x - \underline{x})$$

where $\delta(\cdot)$ denotes the Dirac delta function. With parameter values such that $\gamma > g^Q$, the solution to this ODE is:

$$f^{\mathtt{N}}(x) = \lambda_{\mathtt{N}} \exp[-\lambda_{\mathtt{N}}(x - \underline{x})] \quad \text{where} \quad \lambda_{\mathtt{N}} \equiv \frac{n+\epsilon}{\gamma - g^Q}. \tag{A.12}$$

Similarly, the stationary FP equation for this distribution of log relative quality among old firms is given by:

$$g^{Q}\partial_x f^{\mathtt{O}}(x) - nf^{\mathtt{O}}(x) = -\frac{\epsilon(n+d)}{\epsilon - d} \cdot f^{\mathtt{N}}(x).$$

Dividing through by g^{Q}, and multiplying by the integration factor $\exp(-\lambda_{\mathtt{O}}x)$ where $\lambda_{\mathtt{O}} \equiv n/g^{Q}$, we can rewrite:

$$\partial_x[\exp(-\lambda_{\mathtt{O}}x)f^{\mathtt{O}}(x)] = -\frac{\epsilon(n+d)}{(\epsilon-d)g^{Q}} \cdot \exp(-\lambda_{\mathtt{O}}x)f^{\mathtt{N}}(x).$$

Integrating this equation and solving for $f^{\mathtt{O}}(x)$, we obtain:

$$f^{\mathtt{O}}(x) = \left[\mathcal{C} - \frac{\epsilon(n+d)}{(\epsilon-d)g^{Q}} \cdot \int_{\underline{x}}^{x} f^{\mathtt{N}}(x)\exp(-\lambda_{\mathtt{O}}x)\mathrm{d}x\right]\exp(\lambda_{\mathtt{O}}x) \tag{A.13}$$

where $\mathcal{C}$ is an integration constant. For $f^{\mathtt{O}}(x)$ to be integrable, we must have:

$$\mathcal{C} = \frac{\epsilon(n+d)}{(\epsilon-d)g^{Q}} \cdot \int_{\underline{x}}^{\infty} f^{\mathtt{N}}(x)\exp(-\lambda_{\mathtt{O}}x)\mathrm{d}x.$$

Substituting this expression back into equation (A.13), we obtain:

$$f^{\mathtt{O}}(x) = \frac{\epsilon(n+d)f^{\mathtt{N}}(x)}{(\epsilon-d)(\lambda_{\mathtt{N}}g^{Q}+n)}. \tag{A.14}$$

For $f^{\mathtt{O}}(x)$ to be a probability distribution, we must verify that:

$$\epsilon(n+d) = (\epsilon-d)(\lambda_{\mathtt{N}}g^{Q}+n).$$

Using the consistency condition of the exit rate:

$$d = g^{Q}\mathcal{S}^{\mathtt{O}}\lim_{x\to\underline{x}} f^{\mathtt{O}}(x) \quad \Longrightarrow \quad d = \epsilon g^{Q}/\gamma,$$

it is straightforward to verify that this condition is satisfied. Applying the BGP conditions to equation (A.8), we find:

$$\mathcal{V}^{\mathtt{O}} = \frac{1-\tau^{\mathtt{C}}}{\theta\mathcal{M}[\rho+(\theta-1)g^{Q}]}.$$

To obtain an expression for the growth rate of the average quality index, we combine equations (A.10), (A.11), (A.12), and (A.14), and perform the integration to obtain:

$$g^{Q} = \gamma - \frac{(n+\epsilon)(1-\underline{q}^{\theta-1})}{\theta-1}.$$

Applying the BGP conditions to equation (A.9), we find:

$$\mathcal{V}^{\mathtt{N}} = \frac{\frac{\epsilon+\rho+(\theta-1)g^{Q}}{\rho+(\theta-1)g^{Q}} \cdot \frac{1-\tau^{\mathtt{C}}}{\theta\mathcal{M}}}{\rho+\epsilon-(n+\epsilon)(1-\underline{q}^{\theta-1})+(\theta-1)\gamma/(1+\zeta)}.$$

Finally, applying the BGP conditions to the labor market clearing condition, we obtain:

$$1 = \frac{(1-\tau^{\mathtt{R}})c_{\mathtt{R}}}{\mu\mathcal{V}^{\mathtt{N}}\underline{q}^{\theta-1}} + \frac{(n+d)c_{\mathtt{D}}\gamma^{1+\zeta}}{(1+\zeta)(n+\epsilon)} \cdot \mathcal{M} + c_{\mathtt{R}}(n+d)\mathcal{M}.$$

A.2 The constrained-optimal allocation

In this section, we study the problem of a planner who maximizes utility subject to the same technological constraints as the market, but who does not internalize incumbent-to-incumbent and incumbent-to-entrant technology spillovers. Concretely, the planner chooses research and development allocations taking the aggregate quality index Q_t as *given*. Since Q_t enters both the research and development technologies, treating Q_t as given is equivalent to assuming that the planner does not internalize technology spillovers. In that sense, we refer to the solution to this problem as *constrained*-optimal.

Notation. We introduce the following notation to define the inner product between two square-integrable functions $f(x), g(x) : \Omega \to \mathbb{R}$ over their common domain:

$$\langle f(x), g(x) \rangle_{x\in\Omega} \equiv \int_{\Omega} f(x)g(x)\mathrm{d}x.$$

Second, let us denote the (partial) Gateaux derivative of a functional F with respect to the function $f(x)$ in direction $\varrho(x)$ as:

$$\delta F[f(x);\varrho(x)] \equiv \left.\frac{\partial F[f(x)+\varepsilon\cdot\varrho(x),.]}{\partial\varepsilon}\right|_{\varepsilon=0}$$

where the functional F can take additional arguments through the "dot" notation and the "test" function $\varrho(x)$ is assumed to vanish on the boundaries of the relevant integration domain.

The planner's problem. Consider the problem of a planner seeking to maximize the following objective:

$$U_0 = \int_0^\infty e^{-(\rho-n)t} \ln(C_t/N_t)\mathrm{d}t$$

subject to the constraints:[14]

$$\begin{aligned} C_t &= [\textstyle\sum_{\mathtt{S}\in\{\mathtt{N},\mathtt{O}\}} \langle (q l_t^{\mathtt{PS}}(q))^{\frac{\theta-1}{\theta}}, m_t^{\mathtt{S}}(q)\rangle_{q\in[\underline{q}_t,\infty)}]^{\frac{\theta}{\theta-1}}, \\ N_t &\geq \textstyle\sum_{\mathtt{S}\in\{\mathtt{N},\mathtt{O}\}} \langle l_t^{\mathtt{PS}}(q), m_t^{\mathtt{S}}(q)\rangle_{q\in[\underline{q}_t,\infty)} + \langle l_t^{\mathtt{D}}(q), m_t^{\mathtt{N}}(q)\rangle_{q\in[\underline{q}_t,\infty)} + L_t^{\mathtt{R}}, \\ \dot{m}_t^{\mathtt{N}}(q) &= -\partial_q[\gamma_t(q) q m_t^{\mathtt{N}}(q)] - \epsilon m_t^{\mathtt{N}}(q) + \delta(q-\underline{q}_t) L_t^{\mathtt{R}}/c_{\mathtt{R}}, \\ \dot{m}_t^{\mathtt{O}}(q) &= \epsilon m_t^{\mathtt{N}}(q) - \delta(q-\underline{q}_t) g_t^{Q} m_t^{\mathtt{O}}(q) \end{aligned}$$

by choosing $\{\{\{l_t^{\mathtt{PS}}(q)\}_{\mathtt{S}\in\{\mathtt{N},\mathtt{O}\}}, l_t^{\mathtt{D}}(q)\}_{q=\underline{q}_t}^{\infty}, L_t^{\mathtt{R}}\}_{t=0}^{\infty}$. The solution to the planner's problem is "constrained-optimal" in the sense that Q_t is taken as given. Reformulating this problem using the current-value Hamiltonian, we obtain:[15]

$$\begin{aligned} \mathcal{H}_t &= \ln(C_t/N_t) + \nu_t^L[N_t - \textstyle\sum_{\mathtt{S}\in\{\mathtt{N},\mathtt{O}\}} \langle l_t^{\mathtt{PS}}(q), m_t^{\mathtt{S}}(q)\rangle_{q\in[\underline{q}_t,\infty)} - \langle l_t^{\mathtt{D}}(q), m_t^{\mathtt{N}}(q)\rangle_{q\in[\underline{q}_t,\infty)} - L_t^{\mathtt{R}}] \\ &- \langle \nu_t^{\mathtt{N}}(q), \partial_q[\gamma_t(q) q m_t^{\mathtt{N}}(q)]\rangle_{q\in[\underline{q}_t,\infty)} + \epsilon\langle \nu_t^{\mathtt{O}}(q) - \nu_t^{\mathtt{N}}(q), m_t^{\mathtt{N}}(q)\rangle_{q\in[\underline{q}_t,\infty)} \\ &+ \nu_t^{\mathtt{N}}(\underline{q}_t) L_t^{\mathtt{R}}/c_{\mathtt{R}} \end{aligned}$$

where $\{\nu_t^L, \{\nu_t^{\mathtt{N}}(q), \nu_t^{\mathtt{O}}(q)\}_{q=\underline{q}_t}^{\infty}\}_{t=0}^{\infty}$ are the costate functions. Using integration by parts, we can rewrite:

$$\begin{aligned} \mathcal{H}_t &= \ln(C_t/N_t) + \nu_t^L[N_t - \textstyle\sum_{\mathtt{S}\in\{\mathtt{N},\mathtt{O}\}} \langle l_t^{\mathtt{PS}}(q), m_t^{\mathtt{S}}(q)\rangle_{q\in[\underline{q}_t,\infty)} - \langle l_t^{\mathtt{D}}(q), m_t^{\mathtt{N}}(q)\rangle_{q\in[\underline{q}_t,\infty)} - L_t^{\mathtt{R}}] \\ &+ \langle \gamma_t(q) q \partial_q \nu_t^{\mathtt{N}}(q), m_t^{\mathtt{N}}(q)\rangle_{q\in[\underline{q}_t,\infty)} + \epsilon\langle \nu_t^{\mathtt{O}}(q) - \nu_t^{\mathtt{N}}(q), m_t^{\mathtt{N}}(q)\rangle_{q\in[\underline{q}_t,\infty)} \\ &+ \nu_t^{\mathtt{N}}(\underline{q}_t) L_t^{\mathtt{R}}/c_{\mathtt{R}}. \end{aligned}$$

The first-order condition with respect to $l_t^{\mathtt{PS}}(q)$ implies:

$$\delta C_t[l_t^{\mathtt{PS}}(q); \varrho(q)]/C_t = \nu_t^L \langle m_t^{\mathtt{S}}(q), \varrho(q)\rangle_{q\in[\underline{q}_t,\infty)}$$

[14] Here, we substituted the product resource constraints in the labor resource constraint.

[15] The boundary term associated with the exit threshold does not appear in the Hamiltonian because exit triggers liquidation: when a product line hits $\underline{q}_t$, its liquidation value is transferred to households. In the current-value Hamiltonian, the loss of continuation value from the exit flux is exactly offset by this transfer (in utility units), so the corresponding boundary contributions cancel.

where $\delta C_t[l_t^{\mathtt{PS}}(q);\varrho(q)]$ is the Gateaux derivative of C_t with respect to $l_t^{\mathtt{PS}}(q)$ in direction $\varrho(q)$, which is an arbitrary function that vanishes on the boundaries of $[\underline{q}_t,\infty)$:

$$\delta C_t[l_t^{\mathtt{PS}}(q);\varrho(q)] = C_t^{1/\theta}\langle q^{\frac{\theta-1}{\theta}} l_t^{\mathtt{PS}}(q)^{-1/\theta} m_t^{\mathtt{S}}(q),\varrho(q)\rangle_{q\in[\underline{q}_t,\infty)}.$$

Since that first-order condition must hold for any function $\varrho(q)$, we obtain:

$$l_t^{\mathtt{PS}}(q) = (q/C_t)^{\theta-1}\nu_t^{L-\theta},\ \ \forall q\in[\underline{q}_t,\infty).$$

Integrating this expression, we find:

$$\nu_t^L = 1/L_t^{\mathtt{P}}$$

such that we can rewrite:

$$l_t^{\mathtt{PS}}(q) = (q/Q_t)^{\theta-1}L_t^{\mathtt{P}}/M_t,\ \ \forall q\in[\underline{q}_t,\infty).$$

The first-order condition with respect to $l_t^{\mathtt{D}}(q)$ implies:

$$\langle \gamma_t(q)q\partial_q\nu_t^{\mathtt{N}}(q)m_t^{\mathtt{N}}(q)/l_t^{\mathtt{D}}(q),\varrho(q)\rangle_{q\in[\underline{q}_t,\infty)}/(1+\zeta) = \nu_t^L\langle m_t^{\mathtt{N}}(q),\varrho(q)\rangle_{q\in[\underline{q}_t,\infty)}.$$

Since that first-order condition must hold for any function $\varrho(q)$, we obtain:

$$\gamma_t(q) = \left[\frac{q\partial_q\nu_t^{\mathtt{N}}(q)}{\nu_t^L c_{\mathtt{D}}(q/Q_t)^{\theta-1}}\right]^{1/\zeta},\ \ \forall q\in[\underline{q}_t,\infty).$$

The first-order condition with respect to $L_t^{\mathtt{R}}$ implies:

$$\nu_t^L c_{\mathtt{R}} = \nu_t^{\mathtt{N}}(\underline{q}_t). \tag{A.15}$$

The first-order condition with respect to $m_t^{\mathtt{N}}(q)$ implies:

$$\begin{aligned}&\langle(\rho-n)\nu_t^{\mathtt{N}}(q)-\dot\nu_t^{\mathtt{N}}(q),\varrho(q)\rangle_{q\in[\underline{q}_t,\infty)} = \langle(q/Q_t)^{\theta-1},\varrho(q)\rangle_{q\in[\underline{q}_t,\infty)}/[(\theta-1)M_t]\\&+\langle\gamma_t(q)q\partial_q\nu_t^{\mathtt{N}}(q)-\nu_t^L l_t^{\mathtt{D}}(q)+\epsilon[\nu_t^{\mathtt{O}}(q)-\nu_t^{\mathtt{N}}(q)],\varrho(q)\rangle_{q\in[\underline{q}_t,\infty)}.\end{aligned}$$

Since this condition must hold for any function $\varrho(q)$, we have:

$$\begin{aligned}&(\rho-n)\nu_t^{\mathtt{N}}(q)-\dot\nu_t^{\mathtt{N}}(q) = (q/Q_t)^{\theta-1}/[(\theta-1)M_t]-\nu_t^L l_t^{\mathtt{D}}(q)\\&+\gamma_t(q)q\partial_q\nu_t^{\mathtt{N}}(q)+\epsilon[\nu_t^{\mathtt{O}}(q)-\nu_t^{\mathtt{N}}(q)],\ \ \forall q\in[\underline{q}_t,\infty).\end{aligned} \tag{A.16}$$

Similarly, the first-order condition with respect to $m_t^{\mathtt{O}}(q)$ implies:

$$(\rho - n)\nu_t^{\mathtt{O}}(q) - \dot{\nu}_t^{\mathtt{O}}(q) = (q/Q_t)^{\theta-1}/[(\theta-1)M_t], \ \ \forall q \in (\underline{q}_t, \infty) \tag{A.17}$$

and at the exit threshold, we have:

$$(\rho + g_t^Q - n)\nu_t^{\mathtt{O}}(\underline{q}_t) - \dot{\nu}_t^{\mathtt{O}}(\underline{q}_t) = \underline{q}^{\theta-1}/[(\theta-1)M_t].$$

Defining the following functions:

$$V_t^{\mathtt{N}*}(q) \equiv \nu_t^{\mathtt{N}}(q)C_t/\mu, \ \ V_t^{\mathtt{O}*}(q) \equiv \nu_t^{\mathtt{O}}(q)C_t/\mu, \ \ w_t^* \equiv \nu_t^L C_t/\mu, \ \ r_t^* \equiv \dot{c}_t/c_t + \rho,$$

and substituting them in equation (A.16), we obtain:

$$\begin{aligned} r_t^* V_t^{\mathtt{N}*}(q) - \dot{V}_t^{\mathtt{N}*}(q) = (q/Q_t)^{\theta-1} C_t/(\theta M_t) - w_t^* l_t^{\mathtt{D}}(q) \\ + \gamma_t(q) q \partial_q V_t^{\mathtt{N}*}(q) + \epsilon[V_t^{\mathtt{O}*}(q) - V_t^{\mathtt{N}*}(q)] \end{aligned}$$

where $\gamma_t(q)$ is given by:

$$\gamma_t(q) = \left[\frac{q \partial_q V_t^{\mathtt{N}*}(q)}{c_{\mathtt{D}} w_t^* (q/Q_t)^{\theta-1}} \right]^{1/\zeta}.$$

Doing so for equation (A.17), we obtain:

$$r_t^* V_t^{\mathtt{O}*}(q) - \dot{V}_t^{\mathtt{O}*}(q) = (q/Q_t)^{\theta-1} C_t/(\theta M_t), \ \ \forall [\underline{q}_t, \infty).$$

Finally, substituting these functions in equation (A.15), we obtain:

$$w_t^* c_{\mathtt{R}} = V_t^{\mathtt{N}*}(\underline{q}_t).$$

This demonstrates that market equilibrium allocation is constrained-optimal since it exactly coincides with the solution to the planner's problem, echoing the findings of Dhingra and Morrow (2019). This efficiency holds for both research (entry) and development (product quality improvements). In each case, it arises because two opposing forces exactly cancel each other out under a Dixit and Stiglitz (1977) demand system.[16] For research, a positive consumer surplus externality (entrants do not internalize that they raise the demand for their competitors' products through a love-of-variety) is precisely offset by a negative business-stealing externality (entrants ignore the profits

[16]This knife-edge cancellation follows from not only from Dixit and Stiglitz (1977) demand, but inelastic labor supply, and denominating research and development in units of labor.

they divert from incumbents). Similarly, for development, the incentive to improve a product is dampened by market power, since firms must deploy those quality improvements at a suboptimal production scale. This push toward too little development is perfectly counteracted by the fact that this same underproduction frees up labor for development (away from production), creating a push toward too much. It is worth noting, however, that this efficiency result is sensitive to the model's structure; adding head-to-head creative destruction or a different demand system, for instance, could lead to constrained-inefficient R&D.

By contrast, an *unconstrained* planner—one who recognizes that the allocation of development labor across firms affects the entire product quality distribution and growth—would treat Q_t as an endogenous object. Such a planner could reallocate development labor as to increase the mass of firms near the upper tail of the distribution and thereby raise the average quality index Q_t. Because Q_t enters both the cost of development and the quality at which entrants begin, increasing Q_t lowers future development costs and raises entrants' initial quality. In this sense, the unconstrained planner would explicitly harness incumbent-to-incumbent and incumbent-to-entrant technology spillovers that are taken as given in the constrained problem. Solving this unconstrained planner problem is beyond the scope of this paper.

However, it is also worth emphasizing what we expect the unconstrained planner *not* to do in our environment: This planner would not always find it optimal to shut down research in order to reallocate labor towards development and push product quality growth to its maximal feasible rate of $\overline{\gamma} < \infty$. When quality growth is bounded, diverting labor away from research can raise quality growth only *finitely*, whereas reducing the measure of varieties lowers contemporaneous utility due to a taste for variety. Consequently, collapsing the measure of products cannot deliver the kind of unbounded utility gains emphasized in Trammell (2025).

A related concern is whether an unconstrained planner would want to delete low-quality products in order to improve technology spillovers by mechanically raising Q_t. In our environment, this is not directly feasible: the planner is subject to the same exit technology as the market, so a variety can exit only once its quality falls below the exit threshold $\underline{q}_t \equiv \underline{q} \cdot Q_t$. The planner therefore cannot arbitrarily prune the lower tail of the quality distribution. What the planner *can* do is to reallocate development labor across firms as to increase the mass of products in the upper tail, thereby raising Q_t and the exit threshold with it, as described above.

B Additional figures and tables

This section of the appendix provides additional figures and tables.

Figure B.3: 10% corporate income tax cut

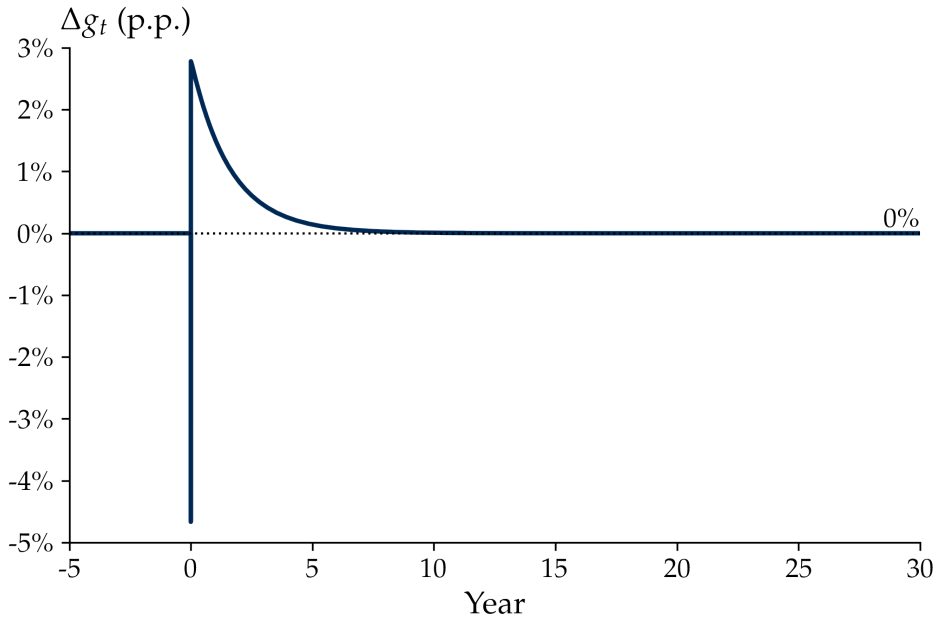

Note: In steady state, this policy has no effect since the equilibrium number of products adjusts to restore average profitability.[17] During the transition, however, higher after-tax profits encourage entry and expand the variety of products available to consumers.

[17]Note, however, that if entry costs were increasing in the number of products per capita (M_t/N_t), such a policy could increase long-run growth by discouraging entry and increasing the return to investments in development. In that sense, a small modification of our model could accommodate long-run market size effects on economic growth.

Figure B.4: Average employment by establishment

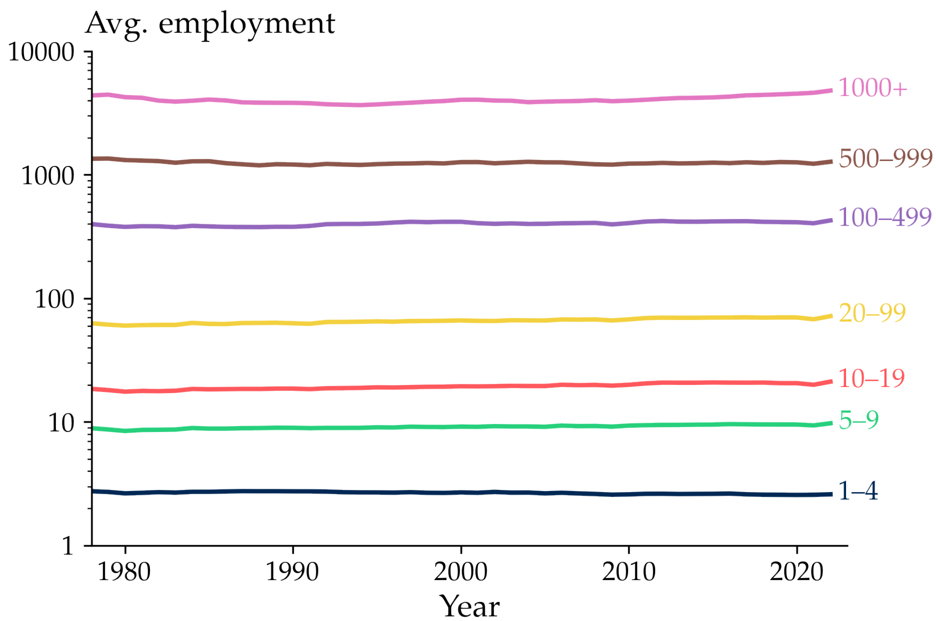

Note: Data from the Business Dynamics Statistics (U.S. Census Bureau, 2023).

Figure B.5: Share of establishments by size class

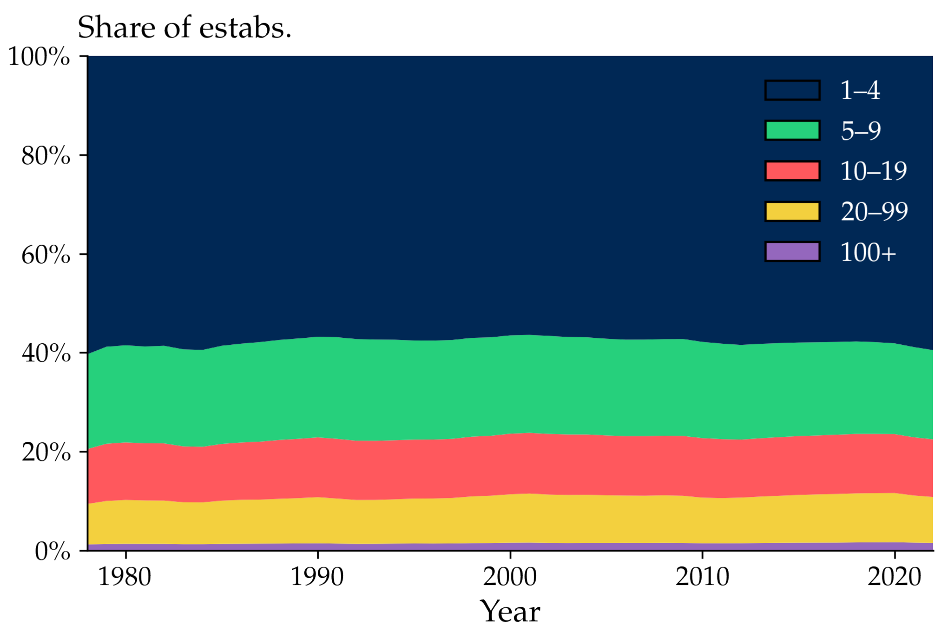

Note: Data from the Business Dynamics Statistics (U.S. Census Bureau, 2023).