Turbulent Growth: Business Dynamism and Aggregate Productivity

Filippo Massari

[a]*Fairfield University, 1073 North Benson Road, Fairfield, 06824, CT, USA*

Abstract

Turbulence is the process of endogenous reallocation of resources (e.g., jobs) across firms due to entry, exit, and churning (movements within the firm-size distribution). This paper formulates a model of turbulent endogenous growth built on the insight that the forces that drive aggregate productivity growth also drive turbulence because the two are manifestations of a single underlying process: profit-driven competition for the market share through innovation. When firms increase their technological knowledge, they gain market share by lowering their price. This reduces the marginal value of further gains in market share. Therefore, the central force in the model is that the incentive to perform cost reducing innovation in house decreases in firms' relative size. Allowing for firm-specific idiosyncratic shocks, this mechanism generates churning and a stationary firm size distribution. These outcomes are robust to introducing entry and exit. A counterfactual exercise studies the effects of a smaller right tail of the R&D productivity distribution of startups. If the decline in job reallocation rates across incumbents observed in the early 2000 is attributable exclusively to that change, the model can explain almost all of the decline in productivity growth in the same period, without invoking a commensurate reduction in R&D.

Keywords: Aggregate Productivity, Firm Dynamics, Turbulence, Endogenous Growth, Firm-Size Distribution.

1. Introduction

Aggregate productivity growth is the key to improving living standards and among the essential factors that economists strive to understand. Modern growth theory and empirics have emphasized R&D as the driver of productivity growth. They point out that firm-level R&D is responsible for firm growth and new firm formation, while lack of it may ultimately lead to firm death. Consequently, aggregate productivity growth is a *turbulent process* intrinsically connected to business dynamism.[1]

In this perspective, the observation that the U.S. (although similar trends are visible in other high-income economies) rates of entry, exit and churning have declined in the past few decades has raised suspicions that the same causes behind this decline could explain the productivity growth slowdown, which occurred around the same time.[2] However, the joint study of aggregate productivity growth and turbulence poses a non-trivial theoretical challenge. Delivering sustained economic growth endogenously requires constant returns to knowledge accumulation. Meanwhile, when allowing for R&D differentials across firms, constant returns facilitate the accumulation of technological knowledge for firms that already possess more. Consequently, this force promotes a dynamic tendency to monopoly, which is counterfactual in most industries. A proper investigation of the phenomenon would have to carefully characterize the process that delivers a stationary firm-size distribution.

This paper develops a theoretical contribution that is necessary to study the association between turbulence and growth. The contribution consists in a novel mechanism that preserves the firm-size distribution stationarity. The mechanism results in the systematic dependence of R&D on market share, thus causing churning. Meanwhile, as R&D is also the driver of long-run growth, the model captures a novel source of interdependence between turbulence and growth. The resulting framework relies on standard industrial organization assumptions and includes a multitude of forces that connect product-level competition and growth. I therefore use it to study the effects of disappearing highly innovative startups on turbulence and aggregate growth by assuming that the right tail of entrants' ability to innovate got thinner.

[1] A turbulent process in this context emerges from repeated changes in the demographics of firms. High entry and exit rates characterize this turbulent business environment. Additionally, most firms expand or contract as they gain or lose market share, a phenomenon known as churning, which entails reshuffling of firms' position within the firm-size distribution. Brown et al. (2008, p. 3) define turbulence to include both these aspects and identify an appropriate measure, namely, the job reallocation rate. They write that turbulence is "the entire process of economic change: worker reallocation as workers change jobs and job reallocation from firms contracting and shutting down, to firms expanding and starting up."

[2] See Naudé (2022) for an example of how economists think of these two phenomena as connected, and for a review of the existing explanations for the observed trends.

The key implication of this structure is that churning is endogenous and a fundamental property of the aggregate growth process. The possibility of stealing market share from others provides the incentive to innovate. However, this force also gives rise to firm-level *economic diminishing returns*, because increasing one's own market share reduces the remaining market share to be gained with further innovation. Consequently, all else constant, incentives to innovate decline with the product market share giving rise to a negative relationship between the firm's growth rate and its relative size.

This mechanism generates mean reversion, therefore delivering a stationary firm-size distribution that features churning, i.e. movements of firms within the distribution. Importantly, the firm-size distribution arises from firms' dynamic optimization and has an endogenous bounded support. Therefore, the shape of the distribution depends on the parameters that regulate firms' incentives to innovate and the speed of mean-reversion. Any parameter change that affects firms' investment decisions would modify the distribution's tails, turbulence, and the economy's growth rate.

Do firm-level economic diminishing returns imply that the aggregate productivity growth rate will tend to zero eventually? No. As these diminishing returns are in *relative* terms, it is possible to accommodate increasing returns to scale in *absolute* terms. The necessary condition for endogenous growth is that returns to innovation are constant or average. Under this condition, R&D investment is constant, and firms grow at a constant rate, implying a constant aggregate growth rate. The implication of firm-level economic diminishing returns is that the distribution of market shares eventually becomes stationary. Therefore, the distribution of returns to innovation is stationary as well, giving rise to a stationary and ergodic distribution of firms' growth rates (with the properties discussed before), and a constant aggregate growth rate. As a result, the aggregate productivity growth rate depends on firms' optimal R&D decisions as in other fully endogenous growth models.

Allowing for simultaneous and endogenous entry and exit does not disrupt this process, while it introduces attractive features. First, it makes the model more realistic by matching a salient feature of the data, i.e. most industries are characterized by high rates of simultaneous entry and exit. Second, it introduces selection effects that determine firms' expected life cycle. Third, it adds a process that affects quantitatively the economy's growth rate, most importantly because the average R&D productivity changes endogenously as a result of these selection effects. Fourth, it affects quantitatively churning because entrants' initial draws of R&D productivity and technological knowledge may be far away from the level at which firms grow at the same rate as the rest of the economy.

The framework I propose builds on the endogenous growth and the firm dynamics literature. A continuum of firms populates the market that features endogenous entry and exit decisions.[3] The crucial theoretical ingredients responsible for the main result are standard: (i) in-house R&D; (ii) a firm-specific idiosyncratic shock drawn from a common distribution, which in this case hits the innovation process that is the root cause of heterogeneity in firms' knowledge stock and market share; and (iii) imperfect substitutability between goods, which produces a demand system that features diminishing returns to relative knowledge — the firm's knowledge stock relative to the average knowledge stock. As firms face a downward-sloping demand curve, expanding their volume of production relative to their competitors requires reducing their price. This exerts downward pressure on the return to further market share expansion, thus reducing the incentive to innovate and grow even larger.

In this paper, I also present a comparative statics exercise. Specifically, as the model allows for dispersion in firms' growth rates, it can study the implications of disappearing high-growth startups, documented for the U.S. in Decker et al. (2016b) around the year 2000. I achieve this goal by reducing the right tail's thickness of entrants' initial draw of ability to innovate. I find that the job reallocation rate across incumbents (a measure of churning) is particularly sensitive to this parameter, and so is the aggregate productivity growth rate, while the entry rate varies little. Changing the relevant parameter to match the decline in job reallocation rate across incumbents before and after 2003 leads to a reduction in growth by 0.7 percentage points, which is almost all of that observed. Interestingly, the reduction in growth is not accompanied by a decline in R&D, which remains unaffected.

Why is churning so sensitive to the presence of high growth startups? The main theoretical theme of this paper provides an explanation. Firms' ability to innovate is ultimately reflected in their size. It is reflected in their growth rate only during their transition from their current size to the size at which they grow at the same rate as the rest of the economy, their no-churning size.

[3] Although this paper's focus is on the product-line as this is the relevant unit for discussing aggregate productivity growth, I use the term firm as a synonymous throughout the paper. Similarly, the word entry refers to horizontal innovation, namely the creation of a new product. Entry is therefore to be conceived as entry of a new product.

Because data suggest that entrants start small, highly innovative startups are quite far from their no-churning size. Hence, they disproportionally account for the level of churning in the economy.

Literature Review. This paper merges the literature on firm dynamics with the endogenous growth literature, specifically Hopenhayn (1992) and Peretto and Connolly (2007), by providing a framework which allows studying entry, exit, churning, the number of goods, the firm-size distribution, and aggregate productivity growth jointly.

As in Hopenhayn (1992), the model delivers endogenously entry, exit, and firm dynamics. These dynamics are driven by idiosyncratic shocks to which firms respond actively by adjusting their size. Contrary to Hopenhayn, where the shock hits firms' productivity, in my model, differences in the productivity level arise endogenously as the outcome of firms' R&D investment. This deviation allows me to derive an endogenous aggregate productivity growth rate jointly determined with the other moments of the model. The framework proposed by Hopenhayn is the foundation of the firm dynamics literature reviewed in Hopenhayn (2014) and in Restuccia and Rogerson (2017), which has recently devoted much attention to resource allocation and the aggregate productivity level. My paper provides a natural extension to this class of models by adding the growth component in a framework that is otherwise the same. In this way, I contribute to the goal that Restuccia and Rogerson (2017, p. 168) identify when discussing future directions for research. They assert: "From a modeling point of view, the key issue is to extend the simple static model of heterogeneous producers [...] to a dynamic setting that includes endogenous decisions that influence future productivity", to "go beyond static effects of misallocation, and focus on the potentially much larger dynamic effects."

Peretto and Connolly (2007), which builds on Peretto (1999), is a useful framework that closely aligns with traditional work in industrial organization. The key attractive feature of this framework, which is absent in other endogenous growth models, is that firms perform R&D in house, thus accounting for the effects of product-level competition at the time of their investment decision. As in Peretto and Connolly, the framework I introduce in this paper includes vertical innovation — cost reduction in the production of existing goods — and horizontal innovation — development of new products — where the former is the engine of long-run growth while the latter drives the equilibrium number of goods, thus the degree of product-level competition. While that model focuses on the symmetric equilibrium, my model introduces an idiosyncratic shock to firms' R&D productivity, thus delivering a non-degenerate firm-size distribution, and richer firm dynamics by giving rise to churning and allowing for simultaneous entry and exit. Introducing these aspects in that model is a particularly significant contribution, considering the proven usefulness of that framework in addressing typical industrial organization questions within a growth context.

Endogenous growth models with heterogeneous producers are not new. The intellectual foundations were laid down in the field of industrial organization, specifically by Ericson and Pakes (1995). When it comes to the interaction between heterogeneous producers, dynamic product-level competition, and growth, the literature of reference consists of two notable papers: Thompson (2001) and Laincz (2009). Contrary to my work, the former paper assumes away the economic diminishing returns to endogenous productivity, thus removing any dependence of R&D on firm-specific market share. In my model, this is an important driver of turbulence, and an element that adds a feedback mechanism from the firm-size distribution to R&D investment, therefore to aggregate productivity growth. Furthermore, my paper introduces exit as an optimal stopping problem, which is absent in Thompson's model. Laincz (2009) delivers a tendency towards monopoly countered only by technological diffusion from the industry leader to entrants. My model, instead, obtains a non-degenerate distribution from assuming product differentiation, and it is therefore complementary to Laincz's model as the two frameworks describe different types of market.

The literature on firm dynamics and growth nowadays builds mostly on the Klette and Kortum (2004) framework.[4] In that framework, firms are conceptually different entities from this paper, and traditional industrial organization. Whereas the Klette and Kortum framework conceives firms as a random collection of products, here firms are organizations that accumulate in house knowledge that enables them to produce a differentiated product. In the Klette and Kortum framework, for mathematical convenience and to avoid the theoretical issue that motivates my paper, R&D is untargeted: firms that innovate successfully improve on a competitor's product without any knowledge of which product it may be at the time of investment. As

[4]Relevant works that build on that framework are Luttmer (2007), Lentz and Mortensen (2008, 2016), Acemoglu et al. (2018), Peters (2020), and De Ridder (2024), which is itself an extension of Akcigit and Kerr (2018) discussed later.

a result, any consideration on market share at the product level is absent. Meanwhile, a general theme in that literature is that specific assumptions, usually on entry and exit, are required for preserving the distribution stationary.[5] The mechanism that I present in this paper makes any of those assumption unnecessary.

The extensions of Klette and Kortum that include own product improvement, thus proposing a model that comes closer to my literature of reference, are two: Acemoglu and Cao (2015) and Akcigit and Kerr (2018).[6] In doing so, they clarify the relevance of the mechanism I show in this paper. In particular, Acemoglu and Cao (2015) explicitly point out that, in their baseline model, allowing for own product improvement leads to a degeneration of the firm size distribution: "the limiting random variable [firm size] only takes either zero or infinity values. In other words, the distribution of normalized firm sizes will continuously expand and a stationary distribution does not exist in a linear BGP [balanced growth path]" (p. 268). To solve this issue, they extend their model by introducing imitative entry: imitators replace incumbents that would have fallen far behind the rest of the market. Importantly, they assume that their initial productivity level is proportional to the average level. As a result, there is always a regeneration of firms that moves the left tail of the distribution such that the right tail never runs away. An analyst that considers entrants' imitation an important feature of the growth process could add it to my model without loss of generality, but this aspect is unnecessary to deliver a non-degenerate firm size distribution.

Similarly, Akcigit and Kerr (2018) requires a specific assumption at the product level to prevent a single product from monopolizing the market. Although one of their main prediction is a decline of firm-level R&D in size, this result does not depend on the dynamics of own product improvement, but stems from a completely different mechanism than the one presented in this paper. In that model, incumbents conduct both horizontal and vertical

[5]For example, many papers either rely on entrants replacing firms that would otherwise become too small, or an exogenous death shock that kills the most successful firm preventing them from becoming too large and monopolizing the market (Klette and Kortum, 2004; Cao et al., 2017; Acemoglu et al., 2018; Akcigit and Ates, 2019).

[6]The idea that productivity growth comes mostly from own product improvement has been well known for a long time, see for example the literature review by Dosi (1988). An empirical exercise documenting that for the U.S. in recent decades is Garcia-Macia et al. (2019).

innovation. The dependence of R&D on size comes exclusively from the horizontal channel. Meanwhile, R&D on vertical innovation is independent of size because the cost of doing R&D is assumed to increase at exactly the same rate as the return from further investment. I relax this knife-edge assumption showing that doing so introduces interdependence between growth and churning, and a market-driven force that counters a tendency towards monopoly. The mechanism for the negative relationship between R&D and firm size described here is therefore complementary to the one presented there.

The rest of the paper is organized as follows. Section 2 presents the model and derives the aggregate productivity growth rate. Section 3 discusses the model's implications with a fixed number of firms and no entry and exit. Section 4 discusses the process of gradual creative destruction and the firm life cycle when entry and exit are endogenous. Section 5 provides the quantitative exercise. Section 6 concludes.

2. A Model of Turbulent Growth

This section describes the model. Time is discrete. A monopolistically competitive intermediate sector consists of a mass of firms that produce a unique good, sold to a perfectly competitive final sector that assembles them in a final good. The final sector distributes the final good to the representative household for consumption, and to entrepreneurs for setting up new firms. Innovation takes up two different forms: the *technological depth* is augmented by improving production processes for existing goods; the *technological breadth* is expanded through the introduction of new goods. Firms invest in R&D to lower their goods' production cost and face an endogenous exit decision at the end of the period. The presence of firm-level idiosyncratic shock R&D productivity generates heterogeneity in firms' productivity levels. New firms enter the market upon payment of a sunk cost in units of output by introducing a new good. All the aggregate variables evolve deterministically.

Households face a consumption/saving choice, along the lines of Bilbiie et al. (2012). Furthermore, they supply labor inelastically.

2.1. Households

The economy is populated by a representative household of size $L_t = L_0(1+\lambda)^t$, where λ is the population growth rate. The household is endowed with L_t units of labor that it supplies inelastically. It makes decisions on how to allocate its income to consumption goods or saving at each point in time.

The representative household maximizes its lifetime utility function,

$$\max_{\{c_t, s_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t L_t \ln c_t, \tag{1}$$

by choosing the sequence of per capita consumption in the final good, c_t, and their saving in a portfolio of stocks of real value s_{t+1}.

The household derives its income from the per capita real wage w_t, and a return r_t on the portfolio of stocks, while it allocates this income to consumption and saving in the portfolio itself. As in Bilbiie et al. (2012), the portfolio is managed by a risk-neutral manager who operates in a perfectly competitive environment. It includes all firms that populate the economy and new firms, whose entry cost is financed by issuing equity. This implies that the idiosyncratic risk is diversified away, simplifying the problem. After normalizing the price index to 1, the household faces the following budget constraint expressed in real terms:

$$s_t + c_t L_t \leq (1 + r_t) s_{t-1} + w_t L_t. \tag{2}$$

Combining the first-order conditions, I obtain the Euler equation that governs the household's saving decision,

$$1 + r_{t+1} = \frac{1}{\beta} \left(\frac{c_{t+1}}{c_t} \right). \tag{3}$$

2.2. Final Sector

I now turn to the description of the economy's production side, starting from the final sector to derive the demand for each intermediate good. A perfectly competitive final sector sells the final good to the household and to entrepreneurs who need it to finance the sunk entry cost. It assembles the final good according to a CES aggregator:

$$Y_t = \left[\int_0^{N_t} x_{i,t}^{\frac{\epsilon-1}{\epsilon}} di \right]^{\frac{\epsilon}{\epsilon-1}}, \tag{4}$$

given a real output Y_t, made from units of the different intermediate goods $x_{i,t}$, the only inputs. N_t is the mass of goods, and $\epsilon > 1$ is the elasticity of substitution across them. The price index, which is chosen as the numeraire, is:

$$P_t = \left[\int_0^{N_t} P_{i,t}^{1-\epsilon} di \right]^{\frac{1}{1-\epsilon}}, \tag{5}$$

where $P_{i,t}$ is the price of each good i.

The representative retailer maximizes his profits by supplying the household and potential entrants with units of the basket of goods. The profit maximization yields the following demand schedule for good i:

$$x_{i,t} = Y_t p_{i,t}^{-\epsilon}, \tag{6}$$

where $p_{i,t} = \frac{P_{i,t}}{P_t}$ is the relative price.

2.3. Intermediate Sector: Production, Innovation, Entry, and Exit

This subsection describes the intermediate sector of the economy. It consists of incumbents, entrants, and exiting firms.

At time t, the intermediate sector is populated by N_t firms with market power producing their own unique good.

The demand schedule derived above implies a revenue per good of:

$$\underbrace{P_{i,t} x_{i,t}}_{\text{Revenue}} = \underbrace{P_t Y_t}_{\substack{\text{Market}\\\text{size}}} \underbrace{p_{i,t}^{1-\epsilon}}_{\substack{\text{Market}\\\text{share}}}, \tag{7}$$

which can be decomposed into market size and market share.[7] The decomposition provides an insight into the competitive process underlying the model. As $\epsilon > 1$, firms can gain market share at others' expense by lowering their relative price. Additionally, two opposite forces affect revenue per good: changes in aggregate spending, namely market size, and changes in the number of producers, which dilute market shares. Market size is beyond the control of the firm, therefore the only way to increase their revenue is for the firm to reduce price and steal market share from others.

The following subsections describe, in turn, the decisions of incumbents and entrants.

[7]By rearranging equation (7) to isolate $p_{i,t}^{1-\epsilon}$, one can observe that it equals the ratio of expenditure on good i and total expenditure, the definition of market share.

2.3.1. Incumbents

Incumbents face a demand given by equation (6). They employ labor that is allocated to produce the intermediate good, $l_{x_{i,t}}$, to cover the fixed costs of production Φ, and to produce knowledge that reduces the future cost of production, namely to perform R&D, $l_{Z_{i,t}}$. They maximize their value, which is the present value of the stream of dividends, by choosing the optimal price, production labor, R&D labor, and whether to exit the market or not.

For the sake of exposition, I break their optimization problem into a static and a dynamic component to derive a cleaner Bellman equation as in other related works, such as Acemoglu et al. (2018). The static component is a per-period dividend maximization, holding constant R&D investment. This allows me to derive an optimal operating profit, conditional on the state, that can be plugged into the Bellman equation. The dynamic component involves an investment decision to maximize the firm's value, with an option to exit the market if it turns negative.

This structure follows Peretto and Connolly (2007), except for the R&D productivity, which is stochastic. This idiosyncratic uncertainty generates heterogeneity in the knowledge stock across firms, and, consequently, in their productivity level.

Timing of Events

The timing of the events is the following: first, an incumbent firm observes its draw for the R&D productivity. Second, it hires labor to produce, invest in R&D, and cover its fixed cost; it picks the price and sells its good. After that, it distributes dividends to the household.[8] Finally, at the end of the period, it decides to exit if its continuation value is negative.

Static Problem: Dividends

In each period, dividends are given by

$$\pi_{i,t} = p_{i,t}x_{i,t} - w_t(l_{x_{i,t}} + l_{Z_{i,t}} + \Phi). \tag{8}$$

where $l_{x_{i,t}}$ is labor allocated to production, and Φ is overhead labor. As mentioned above, to simplify the exposition, I break down the firm maximization problem into a static and a dynamic component. This is equivalent to firms

[8] As there are no liquidity constraints, a negative dividend would imply borrowing from the household through the financial intermediary.

maximizing dividends in each period by choosing how much to produce and what price to charge, holding for the moment R&D labor, $l_{Z_{i,t}}$, constant.

Following the literature, the production technology includes only productivity and labor, such that:

$$x_{i,t} = Z_{i,t}^{\theta} l_{x_{i,t}} \qquad 0 < \theta, \tag{9}$$

where $Z_{i,t}$ is the endogenous stock of knowledge possessed by firm i, and the parameter θ determines the returns to knowledge, or the extent to which production is knowledge-intensive.

The static maximization problem requires a choice of production labor, a price and a quantity to maximize equation (8), subject to demand (6), and the production function (9). The first order conditions yield a production labor demand of:

$$l_{x_{i,t}} = \left[\frac{\epsilon - 1}{\epsilon w_t} \left(\frac{Y_t}{N_t} \right)^{\frac{1}{\epsilon}} Z_{i,t}^{\theta \frac{\epsilon - 1}{\epsilon}} \right]^{\epsilon}. \tag{10}$$

Firms' production labor demand is increasing in the productivity level, Z_t^{θ}, and decreasing in wage. It increases with the overall spending on final goods and declines in the number of goods. In other words, it increases in market share.

A labor demand schedule as in (10) implies a price of:

$$p_{i,t} = \frac{\epsilon}{\epsilon - 1} \frac{w_t}{Z_{i,t}^{\theta}}. \tag{11}$$

This optimal pricing strategy involves charging a constant markup over marginal cost. Importantly, firms can reduce their relative price by improving their technological knowledge.

As anticipated earlier, from equation (7), reducing the relative price, thus gaining market share, is the only way for firms to increase their revenue. Therefore, equation (11) illustrates the fundamental way in which competition occurs: by accumulating technological knowledge faster than the rate of wage growth, firms can lower their relative price and steal market share from competitors. In other words, firms have an incentive to innovate because they can gain market share at the expense of others, and increase their revenue as a result.

Substituting (10) and (11) into equation (8) and using equation (9), dividends can be re-expressed as a function of $Z_{i,t}$, and $l_{Z_{i,t}}$ only.

Heterogeneity and Dynamics: Firm Value Maximization and Exit Decision

Here, I present the dynamic problem of the firm. Each firm makes an investment decision to increase their future knowledge, thus reducing their production cost. The R&D productivity is subject to an idiosyncratic shock, driving heterogeneity in firms' productivity level.

Firms increase their future stock of knowledge through R&D investment. Following Peretto and Smulders (2002), the R&D technology is:

$$Z_{i,t} - Z_{i,t-1} = \alpha_{i,t-1} Z_{i,t-1}^{\mu} Z_{t-1}^{1-\mu} l_{Z_{i,t-1}}^{\zeta}, \tag{12}$$

$0 < \mu < 1$ is a parameter that regulates the private and social returns to knowledge, $\alpha_{i,t} > 0$ is the firm-specific productivity of R&D, and $Z_t \equiv \frac{1}{N_t}\int_0^{N_t} Z_{i,t} di$ is the knowledge spillover, namely, the element that captures the partial non-excludability of knowledge, and the consequent ability of firms to make use of knowledge acquired by others.

An R&D technology of this kind captures four important elements. First, new knowledge is a function of the existing stock of knowledge due to its cumulative nature, i.e. new knowledge builds on existing knowledge. The linearity is the simplest and most tractable specification in which knowledge is the factor that drives long-run exponential growth at a constant rate.[9]

Second, only the firm that produces good i possesses the expertise to improve that line of production, based on the idea that a large driver of innovation is firm-specific in-house technology, widely documented empirically (Dosi, 1988; Garcia-Macia et al., 2019). Therefore, R&D is performed in-house, implying that firms know what product-line they are improving upon when the investment decision is taken. The presence of this element is necessary to deliver the scale dependence in growth rates that generates the mean reversion that produces churning and makes the firm-size distribution stationary.

Third, in line with empirical evidence (Laincz and Peretto, 2006; Ha and Howitt, 2007; Madsen, 2008), the spillover occurs from average knowledge and does not increase with the number of different goods produced in the economy. This specification incorporates the idea that the technological distance between lines of research increases as the product market grows larger, thus diluting away the knowledge spillover and eliminating the scale effect,

[9]Peretto (2018) provides a generalization that allows for new knowledge to exhibit increasing returns in the existing stock of knowledge.

as shown in previous works (Peretto, 1998; Dinopoulos and Thompson, 1998; Young, 1998).

Fourth, the firm-specific shock $\alpha_{i,t}$ follows an AR(1) process:

$$\log \alpha_{i,t} = (1-\rho)\log \overline{\alpha} + \rho \log \alpha_{i,t-1} + \xi_{i,t}, \qquad \begin{aligned} \xi_{i,t} &\sim N(0, \sigma_\xi) \\ 0 &\leq \rho < 1 \end{aligned} \tag{13}$$

where $\xi_{i,t}$ is the draw and σ_ξ its standard deviation.

At the beginning of each period, after observing the draw, each firm invests to maximize its value:

$$\max_{\{l_{Z_{i,t+h}}, Z_{i,t+1+h}\}_{h=0}^{\infty}} V_{i,t} = \pi_{i,t}(l_{Z_{i,t}}, Z_{i,t}) + \max\left\{ \mathbb{E}_t \sum_{h=1}^{\infty} \prod_{q=1}^{h} \frac{1}{1+r_{t+q}} \pi_{i,t+h}(l_{Z_{i,t+1}}, Z_{i,t+1}), 0 \right\} \tag{14}$$

Future profits are discounted using the risk-free interest rate r, which is determined by the representative household's time preferences.

The dynamic optimization can be re-expressed as a Bellman equation:

$$V(Z_{i,t}, \alpha_{i,t}) = \max_{\{l_{Z_{i,t}}\}} \left\{ \pi_{i,t}(Z_{i,t}, l_{Z_{i,t}}) + \frac{1}{1+r_{t+1}} \max\{\mathbb{E}_t V(Z_{i,t+1}, \alpha_{i,t+1}), 0\} \right\} \tag{15}$$

constrained by the knowledge accumulation equation (12).

Lastly, as captured by the max operator, firms face an exit decision at the end of the period. If their value falls below zero, they will decide to dismantle the firm and exit the market permanently. Note that this model does not require any other source of exit, such as a death shock. Exiting the market is fully within the control of the firm.

2.3.2. Entry: Creation of New Goods

I now turn to the description of the entry decision. Entry occurs as long as the present value of the expected stream of dividends exceeds the sunk cost of setting up a firm.

Entrants issue equity to finance the cost of entry. The payment of the sunk cost is in units of output.

When taking the entry decision, entrepreneurs know that their knowledge level in the following period will be drawn out of a lognormal distribution (as the firm size distribution is skewed in the data) around the average knowledge level in the economy. Entry knowledge is given by:

$$Z_{t+1}^d \sim Lognormal(\chi_Z, \sigma_Z^E) Z_{t+1}. \tag{16}$$

Entrepreneurs will also draw their initial R&D productivity from:

$$\alpha^d_{t+1} \sim Lognormal(\chi_\alpha, \sigma^E_\alpha). \tag{17}$$

The distribution is lognormal as entrants display skewness in their employment growth rates (Decker et al., 2016a).

As all potential entrants draw from the same distributions, their value is the same. Given an expected initial level of productivity and productivity of R&D, there is entry at time t as long as:

$$v^E_t = \mathbb{E}_t V(\alpha^d_{t+1}, Z^d_{t+1}) \geq Z^\theta_t N_t^{\frac{1}{\epsilon-1}} f_E, \tag{18}$$

where the right side of the inequality is the entry cost made up of a fixed component f_E and of the technological depth, Z^θ_t, and breadth, $N_t^{\frac{1}{\epsilon-1}}$, of the economy. This specification has a practical purpose: in a growing economy, the entry cost must scale with everything else. Otherwise, as the economy grows richer, setting up new firms would become cheaper, introducing a trend in the entry rate which would be counterfactual. The specification presented here is the simplest one consistent with this property, but not the only one that can deliver it. The idea captured by this specification is that with an increase in the sophistication of the production techniques and of the variety of goods available, the capital required to set up a firm increases proportionally.

Firms set up at time t face the same problem as incumbents in period $t+1$.

2.4. Equilibrium

The equilibrium of the model is defined by:

- a wage w_t, interest rate r_t and price index (5) that firms and the household take as given;
- a demand function (6) from the final sector for the intermediate goods;
- a labor supply L_t, and a demand function for production, overhead and R&D labor;
- an Euler equation (3) for the representative household;
- the free entry condition (18);

- a law of motion of firms:

$$N_{t+1} = N_t + N_{E_t} - N_{X_t}, \tag{19}$$

with N_{E_t} being the mass of entering firms, and N_{X_t} the mass of exiting firms;

- a value function $V(Z_{i,t})$;
- and a distribution $\Gamma_t(z_t)$ of relative knowledge, $z_{i,t}$, where $z_{i,t} = \frac{Z_{i,t}}{Z_t}$,

such that the following conditions hold.

First, the interest rate adjusts to guarantee that the value of the portfolio held by the household equals the aggregation of the value of all firms:

$$s_t = \int_0^{N_t - N_{X_t}} v_{i,t} di + \int_0^{N_{E_t}} v_{i,t}^E di. \tag{20}$$

Second, exploiting equation (5), the prices for each variety are such that they guarantee goods market-clearing:

$$Y_t = c_t L_t + Z_t^\theta N_t^{\frac{1}{\epsilon - 1}} f_E N_{E_t}. \tag{21}$$

Third, the wage adjusts to ensure that quantity of labor demanded by each firm for each activity equals its inelastic supply:

$$L_t = \int_0^{N_t} (l_{x_{i,t}} + l_{Z_{i,t}}) di + \Phi N_t. \tag{22}$$

2.4.1. Steady-State

I solve the model for the stationary steady state equilibrium numerically. I present the stationarized version in Appendix A. Appendix B includes a description of the algorithm used to solve the model.

There exists a time-invariant distribution of firms over relative knowledge $\Gamma(z)$ in steady state that is unique given any initial distribution. The following section describes the forces that make this distribution unique and stationary.

Before introducing output growth, it is useful to define relative (to the arithmetic average) knowledge

$$z_{i,t} = \frac{Z_{i,t}}{Z_t}, \tag{23}$$

and the number of firms per capita

$$n_t = \frac{N_t}{L_t}, \tag{24}$$

which is also the inverse of average firm size and remains constant in steady state, where $N_t/N_{t-1} = 1 + \lambda$. For this class of models, the stationarity of average firm size is discussed in Peretto and Connolly (2007). The basic insight is that as population increases, the market size gets higher, thus increasing operating profits. Larger profits stimulate entry. Entry drags down the average market share, restoring the orginial profit level at the original average firm size.

Furthermore, to simplify the notation, define productivity as:

$$A_{i,t} = Z_{i,t}^{\theta}. \tag{25}$$

Aggregate Productivity Level

From the CES aggregator given by equation (4) and the production function in equation (9), I can express real output per capita as:

$$\frac{Y_t}{L_t} = \underbrace{\underbrace{N_t^{\frac{1}{\epsilon-1}}}_{\substack{\text{Tech.}\\\text{breadth}}} \quad \underbrace{A_t}_{\substack{\text{Tech.}\\\text{depth}}} \quad \underbrace{\left[\frac{1}{N_t}\int_0^{N_t} \left(S_{i,t} N_t \, a_{i,t}\right)^{\frac{\epsilon-1}{\epsilon}} di\right]^{\frac{\epsilon}{\epsilon-1}}}_{\substack{\text{Allocative}\\\text{efficiency}}}}_{\text{Aggregate productivity}} \qquad \underbrace{\frac{L_{x_t}}{L_t}}_{\text{Production effort}} \qquad , \tag{26}$$

where $S_{i,t} = l_{x_{i,t}}/L_{x_t}$ is the production labor share for firm i.[10] Output per capita can be thought of as a combination of productivity and resources devoted to production. The latter element depends on the total fraction of labor devoted to producing units of the intermediate goods — where L_{x_t} denotes aggregate production labor. I focus on this model-consistent definition of aggregate productivity because, as the ultimate interest of any analysis on economic growth is the increase in household's utility, the relevant unit to

[10] I obtain equation (26) by plugging the production function (9) into the CES aggregator (4), then multiplying and dividing by $L_{x_t}A_t$. A_t at the denominator moves into the parenthesis to divide $A_{i,t}$. L_{x_t} moves into the parenthesis to divide $l_{x_{i,t}}$. Next, to isolate the relevance of N_t, I multiply and divide by N_t the factors in the parenthesis. I bring the denominator out, and break it into two parts: one that remains within the square bracket, and the other that comes out denoting the technological breadth (love of variety effect).

consider is the output of the final good — which partly goes to consumption — and the effort exerted to produce it. Other definitions of productivity correlate with this one.

Importantly, all terms that make up the aggregate productivity level are endogenous and depend exclusively on a vector of relative knowledge levels z_t and R&D productivity α_t.

A contribution of this paper is to decompose the aggregate productivity level into various factors that can be linked to firm-level productivity. Due to non-linearities, the dispersion in productivity levels and firm sizes is manifested in the aggregate productivity level. This decomposition shows explicitly how. Aggregate productivity includes three different elements. $N_t^{\frac{1}{\epsilon-1}}$ is the love of variety effect implied by the CES aggregator. This arises out of product differentiation and a preference structure that rewards a larger variety of goods in the market.

The second term describes the technological depth of the economy, and it corresponds to the average productivity across firms, defined as:

$$A_t = \int_0^{N_t} S_{i,t} A_{i,t} di, \tag{27}$$

This model-consistent definition of average productivity is useful because it is also commonly adopted in empirical studies (Foster et al., 2001; Melitz and Polanec, 2015). While in those papers the choice is arbitrary, this model offers a theoretical justification for it.

Finally, the last term shows that aggregate productivity depends on the distribution of weighted firm relative productivities, as:

$$a_{i,t} = \frac{A_{i,t}}{A_t}. \tag{28}$$

This element describes the allocative efficiency of the economy. The distribution of individual productivities matters for aggregate productivity because the CES aggregator is a power mean, which is altered by the firms' relative productivity distribution. To understand why this term is tied to the distribution of productivities and labor share, it is useful to notice that equation (27) can be re-expressed as:

$$\frac{1}{N_t} \int_0^{N_t} S_{i,t} N_t a_{i,t} di = 1. \tag{29}$$

The term labeled *allocative efficiency* above would therefore equal 1 under a symmetric equilibrium, or in a model with additive aggregation of goods. It follows that the term in bracket in equation (26) shows the contribution of the higher moments of the productivity and firm-size distributions to the aggregate productivity level.

Aggregate Productivity Growth Rate

I now shift the focus to the growth rate of aggregate productivity, which, together with the population growth rate, determines the growth rate of output per capita in the long-run.

Proposition 1. *Under a time-invariant distribution of relative productivity levels, the long-run growth rate of aggregate productivity is a function of population growth and of the growth rate of the arithmetic average of firms' productivities.*

Proposition 1 highlights the sources of long-run steady-state growth. Its dependence only on the first moment of the aggregate productivity distribution ensures that firm-level productivity changes are the only relevant factors to consider in steady state. As long as the focus is on the steady state where the firm-size distribution is time-invariant, there is no concern over the aggregation of firm productivity increases. I describe the economic mechanism that delivers the time-invariant distribution in the next section.

The proposition can be expressed in a mathematical form starting from equation (26):

$$1 + g^{productivity} = \underbrace{\left[\frac{n_t(z_t, \alpha_t)}{n_{t-1}(z_{t-1}, \alpha_{t-1})}(1+\lambda)\right]^{\frac{1}{\epsilon-1}}}_{\text{semi-endogenous}} \underbrace{(1 + g_t^A(z_{t-1}, \alpha_{t-1}))}_{\text{average productivity}} \underbrace{\left\{\frac{\frac{1}{N_t}\int_0^{N_t} (S_{i,t}N_t\ a_{i,t})^{\frac{\epsilon-1}{\epsilon}}\, di}{\frac{1}{N_{t-1}}\int_0^{N_{t-1}} (S_{i,t-1}N_{t-1}\ a_{i,t-1})^{\frac{\epsilon-1}{\epsilon}}\, di}\right\}^{\frac{\epsilon}{\epsilon-1}}}_{\text{change in the distribution}}, \quad (30)$$

This expression resembles the one in Peretto and Connolly (2007), with the addition of the last term, which depends on heterogeneity in productivity levels and labor shares. The semi-endogenous component depends only on population growth in steady state as the average firm size is stationary. This term is sometimes referred to as *expanding variety*, and it emerges from the CES aggregator, which rewards a higher number of goods. $g_t^A \equiv A_t/A_{t-1} - 1$ is the growth rate of average productivity between $t-1$ and t, and it will be the focus of the remainder of the paper. Finally, the last term signals that aggregate productivity growth is dependent on changes in the distribution of relative productivity. Nevertheless, given a time-invariant distribution in steady state, the long-run growth rate of aggregate productivity is determined exclusively by the first two terms, while the last one is relevant along the transition, an exploration left for future research.

3. Sources of Firm Growth, Churning, and Stationarity

In contrast with the deterministic and symmetric model proposed by Peretto and Connolly (2007), this model introduces a mean preserving spread to the ability to innovate. By comparing this model to the one where the firm size distribution collapses to a single point, one can study the role that higher-order moments of interest play in shaping the aggregate productivity growth process. However, caution is required as heterogeneity in firm sizes and growth rates is endogenous and interdependent with firms' investment decisions. Therefore, the analysis must proceed with the understanding that the shape of the firm size distribution and all dimensions of turbulence are not exogenous factors that the analyst can arbitrarily change to derive their effect on aggregate variables. They are, instead, outcomes of the same forces that drive economic growth.

The first implication of this interdependence concerns the coexistence of all these phenomena. How can there be (i) growth rate differentials, (ii) constant returns to the growth driving factor (and increasing returns to the private factors overall), and (iii) a stationary firm size distribution? The elements (i) and (ii) may suggest a higher growth rate for relatively larger firms, thus promoting a tendency towards monopoly and violating element (iii).

To better illustrate the mechanism, I temporarily shut down entry and exit. I will reintroduce them in the next section where I discuss their role in detail. I do that by setting the parameters $\lambda = 0$, $f_E = \infty$. Notice that I do not need to remove the free exit condition to rule out exit in steady state. When entry is made impossible, any firm that exits the market will increase the profitability of the remaining ones. As this process continues, no firm has a negative continuation value at some point, exit ceases, and the average firm size remains constant. Therefore, the rest of this section will disregard the exit process as it is irrelevant to this section's purpose.

The next two subsections discuss the mechanism that preserves the firm size distribution stationarity and its implications.

3.1. Growth Rate Differentials without Entry and Exit

In what follows, I show the sources of growth rate differentials across firms, which cause churning. Additionally, I discuss the conditions under which growth rates are decreasing in relative knowledge conditional on R&D productivity. This negative relation is what drives the distribution stationarity.

Notably, the model does not necessitate entry and exit to produce this result. This feature contrasts with the majority of growth models with heterogenous firms. The Klette and Kortum (2004, p. 1000) model, on which many recent works are based, requires entry and exit because "without entry, the mass of firms continually declines, the average size of surviving firms becomes ever larger, and the size distribution of survivors becomes ever more skewed." I described similar issues for other models in the literature review.

To show the convergence process to a stationary distribution, I focus on the partial equilibrium of the model. The general equilibrium effects are not fundamentally different from those discussed in Peretto and Connolly (2007).

Using the approximation $g^{A_i}_{t+1} \approx \theta g^{Z_i}_{t+1}$, I can derive the growth rate of an arbitrary firm from equation (12) after plugging in the optimal l_Z value:

$$g^{A_i}_{t+1} \approx \alpha_{i,t+1} z^{\mu-1}_{i,t} l_{Z_{i,t}}(z_{i,t}, \mathbb{E}_t \alpha_{i,t+1})^{\zeta}. \tag{31}$$

This growth rate depends on three elements: the R&D productivity, the initial relative knowledge level, and the R&D effort exerted by the firm.

The term $z^{\mu-1}$ implies that for any given R&D investment and expectation of R&D productivity, the growth rate declines in relative knowledge since private returns to knowledge $\mu < 1$. The knowledge spillover drives this effect by operating as a force of attraction: firms above the average knowledge level will be dragged down in relative terms by the spillover, while firms below the average knowledge level will be lifted by it. The larger the firm, the larger the R&D productivity and investment required to balance this force of attraction. Private returns to knowledge partially offset this effect by facilitating the accumulation of knowledge for firms that already possess more.

Second, R&D is a function of relative knowledge. A necessary condition to deliver this dependence is that the innovator must know in advance what product line they will improve when making the investment decision. However, this condition is absent in several growth models. In this framework, firms perform R&D in-house, in line with empirical evidence (Dosi, 1988). Therefore, when investing, firms consider their relative productivity level — and, consequently, market share.

The value maximization yields the following policy rule:

$$l_{Z_{i,t}}^{1-\zeta} = \frac{1}{1+r_{t+1}} \frac{\alpha_{i,t} z_{i,t}^{\mu}}{w_t \left(1+g_{t+1}^{Z}\right)} \mathbb{E}_t \left[\frac{w_{t+1} l_{Z_{i,t+1}}^{1-\zeta}}{\alpha_{i,t+1} z_{i,t+1}^{\mu}} + \zeta \frac{\partial \pi_{i,t+1}}{\partial z_{i,t+1}} + \frac{\mu w_{t+1} l_{Z_{i,t+1}}}{z_{i,t+1}} \right], \tag{32}$$

with

$$\frac{\partial \pi_{i,t}}{\partial z_{i,t}} = w_t^{1-\epsilon} \left(\frac{\epsilon - 1}{\epsilon} \right)^{\epsilon} \frac{y_t}{n_t} \theta z_{i,t}^{\theta(\epsilon-1)-1}. \tag{33}$$

Equation (32) shows that firms will choose R&D investment by balancing the present value of relative knowledge's marginal benefit and marginal cost. The term outside the bracket is the inverse of the marginal cost of new relative knowledge — with a slight modification as I have kept the diminishing returns to R&D on the left side. It increases with the price of R&D and with average knowledge growth, as faster average knowledge growth requires more investment for firms to keep up with the others. It instead decreases with R&D productivity and relative knowledge to the extent that firms internalize it, as these two elements determine the efficacy of R&D.

The first term in the bracket is the following period's marginal cost of creating new relative knowledge. Firms smooth their R&D investment over time while preferring larger investments in periods when it is cheaper.

The second and third elements in the square bracket are the marginal benefit of creating relative knowledge. The first obvious reason to create new relative knowledge is to increase profits. Additionally, if knowledge creation is facilitated by the internal stock of knowledge within the firm, namely if $\mu > 0$, firms have an extra incentive to invest as their current investment will be beneficial when investing in future periods.

Regarding growth rate differentials and the stationarity of the firm size distribution, the key question regards the relation between R&D investment and relative knowledge.

Profit is concave in relative knowledge as long as $(\epsilon - 1)\theta < 1$. Therefore, the relation between the incentive to innovate and the relative knowledge level of the firm depends on some crucial parameters: μ, ζ, θ, and ϵ, which represent respectively the private returns to knowledge in new knowledge creation, the strength of diminishing returns to R&D, the elasticity of production with respect to the stock of knowledge, and the degree of substitutability across goods which determines the elasticity of demand for each good.

In particular, the ability of firms to internalize the knowledge they produce and exclude others from accessing it constitutes a force of divergence: firms that possess more knowledge are also better able to create more of it, thus reinforcing their advantage over time. Furthermore, the degree of knowledge intensity of the economy compounds this effect by mapping differences in knowledge levels into differences in firm size, production and, ultimately, profits. It is immediately visible that parameters θ and μ are essential in determining the shape of the firm size distribution and raise concerns about its non-degeneracy. The former parameter regulates the degree of increasing returns to the private factors in production. The latter introduces an additional reward to a private factor by facilitating its accumulation.

The other two parameters counter these forces of divergence. Of particular interest is the role of ϵ. Indeed, the condition for concavity of profits in relative knowledge requires either diminishing returns to knowledge in production, or a low enough elasticity of substitution. In the presence of product differentiation, consumers' preference for variety ensures that the most productive good will not be the only one sold. If this preference for variety is strong enough, the incentive for technologically advanced firms to improve their productivity faster than their competitors is overwhelmed by the inability to gain enough market share to justify the effort.

In other words, diminishing returns to relative size originate from the demand side through a mechanism that resembles the one Acemoglu and Ventura (2002) emphasized in a different context. Firms that gain more technological knowledge relative to others increase their production volume. By producing more, they face a lower price as product differentiation ensures that firms face a downward-sloping demand curve. This price reduction is, in turn, responsible for dragging down the return to further knowledge accumulation. As a result, incentives to innovate decline as firms grow larger relative to others.

This mechanism is absent in other models. In fact, the outcomes of existing models in the literature are special cases of this one.

Most models rely on assumptions that deliver firm growth rates consistent with Gibrat Law (Klette and Kortum, 2004; Acemoglu and Cao, 2015; Acemoglu et al., 2018), i.e., the hypothesized absence of any correlation between firm growth and relative size. Gibrat Law is a knife edge case that arises in this model with combinations of parameters such that the forces of divergence are perfectly balanced with the forces of convergence. Although it is a mathematically convenient outcome to match, Gibrat Law is at odds with the data (Sutton, 1997), especially in manufacturing industries (Audretsch et al., 2004), which arguably perform more R&D than service industries. Additionally, as discussed in those papers, delivering a stationary firm-size distribution while matching Gibrat Law requires the presence of entry and exit and specific assumptions on those processes.

Akcigit and Kerr (2018) build on Klette and Kortum (2004), but deviate from Gibrat Law. However, the deviation from Gibrat Law occurs because incumbents perform horizontal innovation, and their innovation intensity decreases with size. Instead, their vertical innovation intensity, which according to empirical evidence it is what matters the most for growth (Garcia-Macia et al., 2019), is independent of relative firm size. Therefore, that model and the one presented here offer complementary explanations for deviations from Gibrat Law.

Another model of interest is the one illustrated in Thompson (2001). In there, growth rates are decreasing in relative size even though R&D is independent of it. My model can reproduce this outcome if $\mu = 0$ and $\theta = 1/(\epsilon - 1)$. The last restriction is the one that would make the firm value linear in relative knowledge (or relative quality in that paper), simplifying the mathematical structure but removing an element of interest.

Finally, the parameters could be calibrated to deliver faster growth for larger firms, thus manifesting a tendency toward monopoly. In this case, the model would deliver an outcome that, in the limit, resembles the one of Aghion and Howitt (1992), as ϵ tends toward infinity. In that model, a monopolist supplies the entire market and has no incentive to innovate. Instead, growth results from innovation introduced by entrants who capture the entire market when they are successful.

Laincz (2009) captures the same idea. In his model, the most productive firm has, on average, a stronger incentive to innovate. Knowledge spillovers from the largest incumbents to entrants counter the tendency toward monopoly. The result is a highly concentrated industry where the largest incumbent keeps innovating to escape the competition of entrants. The model presented here would fall under that realm for parametrizations where the forces of divergence prevail.

What is the right parametrization? Differences in estimates of the Gibrat coefficients, shown for example in Bottazzi et al. (2007), suggest that it will vary by industry. While some industries may manifest a tendency toward monopoly, most industries do not. Consequently, a flexible model delivering any of these possible outcomes is particularly attractive.

In what follows, I focus on the case in which the forces of convergence prevail. The existing literature has addressed all other cases.

3.2. Prediction 1: Stationarity, and Endogenous Churning

This subsection discusses the first relevant set of predictions of the model. Although the production technology exhibits increasing returns to the private factors of production, the firm-size distribution is stationary and non-degenerate for parametrizations that deliver declining growth rates in firm relative sizes. Consequently, churning arises endogenously in the form of conditional mean-reversion. Notably, the strength of this phenomenon depends on the R&D investment decisions of firms.

Figure 1 illustrates the phase diagram that describes this convergence process by showing the expected evolution of firms over relative productivity and R&D productivity. The LL locus shows the long-run expectation of the exogenous AR(1) process that characterizes the evolution of R&D productivity, namely the unconditional expectation of equation (13).

The convergence process over relative productivity can be understood by analyzing the R&D technology given in equation (12), combined with the policy function (32), which can be rearranged to yield:

$$\frac{a_{i,t+1}(z_{i,t+1})}{a_{i,t}(z_{i,t})} = \frac{(1 + \alpha_{i,t} z_{i,t}^{\mu-1} l_{Z_{i,t}}(\alpha_{i,t}, z_{i,t})^{\zeta})^{\theta}}{1 + g_{t+1}^{A}}. \tag{34}$$

At this point, it is possible to construct a locus over $a_{i,t}$ and $\alpha_{i,t}$, along which the relative productivity level remains constant over time. I call this the *no-churning locus*. It is given by:

$$1 + g_{t+1}^{A} = (1 + \alpha_{i,t} z_{i,t}^{\mu-1} l_{Z_{i,t}}(\alpha_{i,t}, z_{i,t})^{\zeta})^{\theta}. \tag{35}$$

This no-churning locus shows the values of relative productivity and R&D productivity at which productivity growth rates equal the average productivity growth rate, which firms take as given.

For any R&D productivity level, convergence to the no-churning locus requires firms' growth rates to decline in relative productivity.

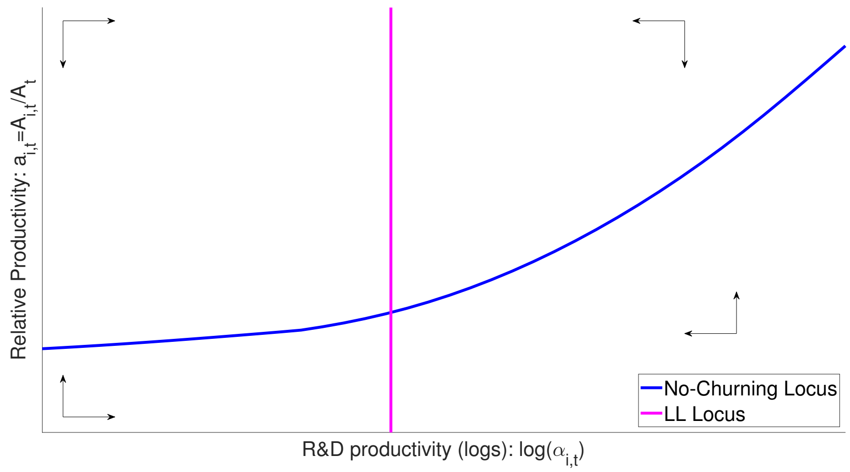

Figure 1: Turbulence, firm growth, and stability.

Note: Phase diagram over the state variable and the exogenous shock. The no-churning locus represents the combination of relative productivity and R&D productivity for which the firm's growth rate equals the growth rate of average productivity. The LL locus shows the long-run expectation of R&D productivity.

As the problem is stochastic, firms are virtually never on the no-churning locus. Therefore, growth rates are not equalized in each period, but only on average. Firms below it will grow faster than average productivity and vice-versa. This mean-reversion process arises due to the forces of attraction discussed above, a combination of scale dependence in R&D investment and knowledge spillovers. Although firms tend endogenously towards the no-churning locus, the shock disrupts their position in the state-space every period. This is one of the key results of the paper: churning, hence turbulence, arises endogenously as the result of firms' optimization.

Furthermore, as Figure 1 shows, the no-churning locus is upward-sloping. This positive slope illustrates that firms with a persistently higher ability to innovate will eventually manifest it in their relative size and not in their growth rate. Understanding why this positive slope arises is crucial to reconcile two seemingly contradictory aspects of the firm growth process. First, firms' productivity growth is strictly increasing in R&D productivity, meaning that more innovative firms grow faster, all else constant. Second, more innovative firms necessitate less investment or knowledge spillovers to maintain their position within the relative productivity distribution. Therefore, how is it possible that the most innovative firm does not grow faster than others forever, thus monopolizing the market in the limit? As more innovative firms grow relatively larger, their growing size is responsible for reducing their investment. Eventually, these firms will reach a level of relative productivity such that the forces of attraction are strong enough to balance their high ability to innovate, thus leading them to grow at the same rate as average productivity.

In addition, the no-churning locus's position is endogenous and depends on the aggregate variables. If the shocks were shut down in steady state the no-churning locus would intersect the LL locus at a relative productivity level of 1. All firms would have the same relative productivity level and grow at the same rate. Systematic churning ensures that the growth rate of a firm with the average relative productivity level and the average R&D productivity level does not grow at the same rate as average productivity. The nonlinearities required for the model to produce endogenous growth and a stationary firm size distribution imply that the first moment of the R&D productivity and relative productivity distributions are not enough to describe the aggregate productivity growth process accurately.

3.3. Stationarity in a Simplified Environment

This section proves stationarity after introducing a few simplifying assumptions for the purpose of mathematical tractability.

Proposition 2. *In partial equilibrium, and under the following simplificatory assumptions:*

- $\alpha_{i,t}$ *evolves deterministically. It is still strictly positive and finite, and its distribution is always stationary;*
- $\mu = 0$*;*
- *Agents are sufficiently impatient to ensure that discounting by two periods yields approximately 0, while discounting by one period yields positive values.*

the firm size distribution is stationary if $\theta(\epsilon - 1) < 1$.

The first assumption helps by removing the expectation operator from the policy function. The simplest process for $\alpha_{i,t}$ conceivable under these assumption is a time invariant one, which implies permanent differences across firms' abilities to innovate. Proving stationarity under fixed differences is particularly noteworthy as it highlights one of the strengths of this model relative to the existing literature: stationarity of the firm size distribution does not arise out of assuming that differences across firms fade away with time. Instead, it arises because of an endogenous market mechanism arising from standard industrial organization assumptions that disciplines firms' optimal investment decisions.

The second assumption, instead, is useful for algebraic convenience, but it is stronger than needed. The same proposition, although with a different and stronger condition for stationarity, could be formulated while relaxing that assumption. The suggestion for analysts who wish to use this model with $\mu > 0$, as I do in my quantitative application, is to check after calibrating it whether the conditions for stationarity are met, i.e. firms' growth rates as a function of relative knowledge or productivity cross the growth rate of average knowledge or productivity only once and from above in the relevant state-space.

The third assumption as well is motivated by algebraic convenience, as it simplifies the differentiation of the policy function with respect to relative knowledge. However, everywhere else, I write the general equations without relying on that assumption.

The proof for proposition 2 is in Appendix C, while here I provide the intuition and the key equation. It relies on demonstrating that (i) knowledge grows at a rate that is strictly decreasing in relative knowledge, and (ii) the range of growth rates as a function of relative knowledge includes the growth rate of average knowledge. These two conditions imply that firms that are below the average level of knowledge will grow faster than average, thus converging to the average knowledge level. Firms that are below it will grow slower, thus converging to average as well.

With these two assumptions, (32) simplifies considerably. First, the expectation operator becomes unnecessary as the problem is now one of perfect foresight. Then the last term in the square bracket drops out. As it is easier to work with stationarized variables, specifically, I define for any variable X_t, $\widehat{X}_t = X_t / Z_t^{\theta} N_t^{\frac{1}{\epsilon - 1}}$ (Appendix A provides all relevant equations with only stationary variables). Iterating forward that equation then leads to:

$$l_{Z_{i,t}} = \left[\sum_{h=1}^{\infty}\prod_{q=1}^{h}\frac{B_{t+1+q}\alpha_{i,t}\zeta^{q}}{\widehat{w}_t}\frac{\partial\widehat{\pi}_{i,t+h}}{\partial z_{i,t+h}}\right]^{\frac{1}{1-\zeta}} + \lim_{s\to\infty}\prod_{q=1}^{\infty}\alpha_{i,t}\widehat{w}_s\frac{B_{t+1+q}}{\alpha_{i,s}\widehat{w}_t} \tag{36}$$

where

$$B_{t+1} = \frac{(1+g^{Z}_{t+1})^{\theta-1}(1+\lambda)^{\frac{1}{\epsilon-1}}\left(\frac{n_{t+1}}{n_t}\right)^{\frac{1}{\epsilon-1}}}{1+r_{t+1}}, \tag{37}$$

and the second term in equation (36) equals 0 because it is an infinite discounting of a stationary term with a discount factor that is smaller than 1.

To ensure that knowledge growth is strictly decreasing in relative knowledge, it is sufficient (although not necessary) to prove that R&D investment is strictly decreasing in relative knowledge, because with $\mu = 0$, equation (31) could depend positively on relative knowledge only through R&D investment. I use implicit differentiation to find $\partial l_{Z_{i,t}}/\partial z_{i,t}$, and verify that it is strictly negative when $\theta(\epsilon-1) < 1$.

Next, I prove condition (ii). Since knowledge growth is strictly decreasing, and growth of average knowledge is strictly positive and finite, a sufficient condition for this proof is that knowledge growth tends to infinity when relative knowledge tends to zero, and it tends to zero when relative knowledge tends to infinity. This part is described in full in the appendix.

4. Entry and Exit

In this section, I relax the assumptions introduced in the previous section to analyze how the process of entry and exit interacts with the rest. Specifically, I remove restrictions on the parameters f_E and λ. However, I still work in an environment where the forces of convergence prevail. A positive fixed cost of production can turn the firm value negative, thus allowing for exit. A finite value for the entry fee can make entry possible. As entry occurs, incumbents face competition from entrants, and the continuation value of some of them becomes negative, thus forcing them to exit. Finally, a positive population growth rate ensures that the steady-state net entry rate is positive, as explained above.

What ensures that the presence of entry and exit will preserve the results illustrated above regarding the stationarity and non-degeneracy of the firm-size distribution? The model has one firm-specific state variable, the knowledge stock, and the firm-specific exogenous shock. The exogenous shock is stationary by assumption. Moreover, with parametrizations considered in the previous section — meaning those that deliver productivity growth rates that decline in relative productivity crossing the average productivity growth rate only once and from above — relative knowledge evolves according to a stationary process. Consequently, the conditions used in Hopenhayn (1992) to prove the existence of a stationary equilibrium with entry and exit are verified. These conditions consist of a stationary process of R&D productivity and relative knowledge; a value function that is strictly increasing and continuous in R&D productivity and relative knowledge and strictly decreasing in the number of firms; and an entry cost below a threshold to allow entry.

It is important to consider a caveat when comparing this model to the one introduced by Hopenhayn. As the growth rate of the economy is positive, the entry cost cannot simply be a parameter. It must instead increase as the economy gains efficiency over time to avoid introducing a trend in entry and exit rates. If technological knowledge increases, affecting relative prices of production or R&D and entry, the resources devoted to each activity will exhibit a different trend, thus introducing a trend in the entry rate. In the model specification presented above, the cost of entry in terms of units of goods increases at the rate of change in the economy's technological breadth and depth. In this way, the entry cost keeps pace with the costs of production and R&D, namely the real wage.

The intuition behind Hopenhayn's proof that is valid here is that an adjustment in the number of firms is the mechanism that balances the entry and exit rates through an effect on the profitability of firms. If the entry rate exceeds the exit rate by more than the population growth rate, the number of firms per capita will rise over time, thus depressing profits as demand spreads over more products. This reduction in profits would lead fewer firms to enter the market and more firms to exit until entry and exit rates are such that the number of firms remains constant over time.

There is, however, a significant conceptual difference relative to Hopenhayn's model. Differences in productivity levels across firms are the endogenous outcome of their investment decisions, as opposed to the outcome of a shock. While this difference does not disrupt the results obtained in Hopenhayn as long as relative productivity evolves according to a stationary process, its relevance is noteworthy first because the productivity distribution across firms is the result of firms' choices; second, because the steady-state aggregate growth rate of the economy is endogenous and dependent on firm-level investment decisions.

The following subsection illustrate the role that entry and exit play in the model. The one after presents some economic implications of adding entry and exit.

4.1. Effects Linking Entry and Growth

The relationship between entry, exit, and growth is complex because multiple effects occur at the same time. In this subsection, I illustrate them.

Effect 1: Expanding Variety

Assuming a CES form for the aggreagation of different goods implies a love for variety. This effect matters for growth as shown in (26), where the output growth rate depends on the technological breadth of production. An increase in the number of products occurs when net entry is positive, implying that an increase in net entry will contribute positively to growth through this channel.

Effect 2: Level Replacement

The model allows for simultaneous entry and exit. In addition, entrants and exiters have on average a different productivity level. The average productivity level of exiters is fully endogenous, whereas the average productivity level of entrants depends on the parameters of equation (16). If entrants have on average a higher productivity level than exiters, more entry and exit will increase growth through this channel, and vice-versa. In a way, this channel has the same aggregate implications for growth of the class of models built on Aghion and Howitt (1992), although the mechanism that leads to replacement is different, and in this model entry does not need to occur in the same industry as the one of exiters. However, an important difference is that in this model this effect will tend to weaken with an increase in exit. That happens because exiters tend to be the least productive firms. Therefore, expanding the exit zone wil lead to more productive firms exiting the market. In the limiting case of 100% exit, the relative productivity level of exiters is 1.

Effect 3: Growth Replacement

This effect is conceptually the same as the previous one, with the difference that the firm-level variable of interest is not productivity, but R&D productivity. If entrants are on average better innovators than exiters, any level of entry rate matched by exit rate will effectively determine a substitution of worse innovators with better innovators, raising the average ability to innovate in the economy. This is the main effect emphasized by the literature on firm dynamics and growth, for example Acemoglu et al. (2018).

Effect 4: Cost Spreading

A parameter change that facilitates entry, will increase the steady state number of firms. Conversely, a parameter change that facilitates exit will decrease the steady state number of firms. A change in the number of firms is a change in the average market share, which has consequences for R&D investment. As visible in equation (32), R&D investment decreases in the number of firms. The cost spreading effect explains this result: as the cost of innovation is spread on all units sold, a more crowded market where each firm sells fewer units reduces the incentives to innovate. This effect is emphasized in the symmetric version of this model (Peretto and Connolly, 2007), but tends to be absent in other firm dynamics models where the number of firms is fixed.

Effect 5: R&D-Growth Decoupling

A final effect relates to the firm-size distribution. As the return to R&D investment depends on the R&D productivity but also on the market share, any force that affects the firm size distribution will affect the returns to R&D in different ways for different firms. Therefore, a change in aggregate R&D does not necessarily mean that each firm changed its R&D in the same way. It could be that firms that are on average worse innovators increased their R&D, while firms that are better innovators decreased it, with an effect on growth that could go in either direction. Because of this effect, the model can conceive a change in aggregate R&D accompanied by a change in the opposite direction in growth. Anything that affects the entry rate will have this effect for two reasons: first, the firm size distribution is an average of the continuing incumbents distribution and the entrants' distribution weighted by the entry rate; second, a change in the entry rate will affect the exit rate, thus modifying the distribution of continuing incumbents.

4.2. Prediction 2: Firm Life Cycle

As in Hopenhayn, the model provides predictions over firms' life cycle. However, as a firm's productivity is not a random draw, the life cycle differs.

Figure 2 re-proposes the phase diagram of the previous section after changing parameters to allow for entry and exit. The first noticeable difference is the presence of an exit locus. This locus is an absorbing barrier: firms whose relative productivity and R&D productivity levels lie below that curve have a negative continuation value and exit at the end of the period. Unlike models of firm dynamics, as the firm's value depends on the endogenous state variable, the distribution over productivity does not have an abrupt truncation but a smoother left tail.

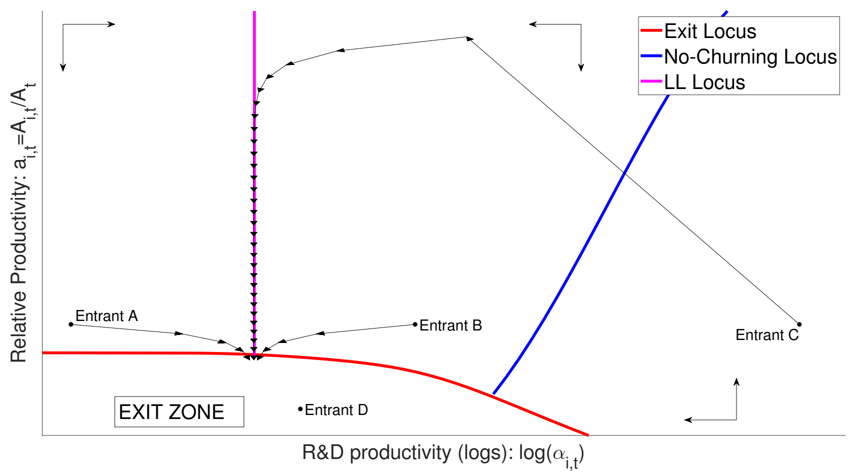

Figure 2: The firm expected life cycle.

Note: Phase diagram over the state variable and the exogenous shock. The no-churning locus represents the combination of relative productivity and R&D productivity for which the firm's growth rate equals the growth rate of average productivity. The LL locus shows the long-run expectation of R&D productivity. The exit locus shows the combination of relative productivity and R&D productivity below which the firm exits the market. The figure also shows the expected life path of four startups that differ in their initial R&D productivity draw.

The first implication of endogenous exit is that R&D productivity's cross-sectional average is higher than its unconditional expectation depicted on the LL locus. This inequality arises because exiting firms have, on average, a lower R&D productivity. Furthermore, with a realistic calibration, entrants have a higher R&D productivity than incumbents. Thus, entry increases the R&D productivity's cross-sectional average even further.

The disconnect between the R&D productivity cross-sectional average and the individual firm's time average implies that each firm has an expected finite life. Figure 2 shows that all firms expect to converge to the exit locus eventually.

These dynamics highlight the relevance of the firm life-cycle for aggregate productivity growth. Surviving entrants with a higher ability to innovate than incumbents gain relative productivity over time, stealing their market share. As a result, less innovative firms lose ground until they exit the market when their relative productivity level is low enough to make them unprofitable.

The phase diagram illustrates this point by including the expected evolution of R&D productivity and relative productivity for four arbitrary entrants. The first three firms differ exclusively in their initial draw of R&D productivity. However, they all enter at the same relative productivity level.

Entrant A is a startup that is subject to a bad draw. Since it enters with little ability to innovate, it will shrink in size and exit after a few periods. Although its ability to innovate will, in expectation, increase over time, reverting to the mean, it will never be sufficiently high for the firm to grow larger.

Entrant B is a startup that enters with a higher ability to innovate. However, its R&D productivity is low enough for it to never increase in relative size. Contrary to firm A, its R&D productivity is expected to shrink over time, but the evolution of relative productivity is the same. Firms A and B represent the median startup in the U.S., which exits the market in less than a decade (based on U.S. Census data).

Entrant C is, instead, what is commonly known as a gazelle, namely a firm that grows at a fast pace. This highly innovative startup type can transform these good ideas into a high productivity growth rate. As the firm improves its relative productivity, it also increases in size and market share, thus reaching the no-churning locus. At that point, its R&D investment level becomes just enough to maintain its size. Meanwhile, as the initial good ideas are explored, and the ability to turn them into new productivity gains fades away, the quality of its new ideas reverts to the mean (in the absence of any new good draw). The firm will, therefore, begin to shrink as other more innovative firms gain market share at its expense. If this firm does not stumble upon any new good idea (i.e., has a good draw), it will eventually become unprofitable and exit the market. This process could be considered a form of *creative destruction*, where the producer of a good (for example, a DVD player producer) drove an imperfect substitute (for example, a VHS player producer) out of the market over time by gradually increasing its relative efficiency.

Entrant D is a startup with bad draws of both relative knowledge and R&D productivity. As a result, it exits the market on the first occasion. In U.S. Census data, these firms are relatively common as about 20% of the population exits the market within their first year of operation.

Since the quantitative exercise proposed in this model concerns the disappearance of high R&D productivity entrants, what does the model say about their effect on churning? Because entrants start small on average, high R&D productivity entrants will cluster in the bottom right corner of Figure 2. Therefore, they will initially be far from the no-churning locus. As a result, they will grow fast to reach it, thus contributing disproportionately to churning. Furthermore, to the extent that they also drive growth of average productivity, they will promote the shrinkage of other firms, adding to churning even more.

5. Quantitative Exercise

This section starts by introducing the calibrated model. The first goal of the calibration is to show that the growth rates of productivity cross the average productivity growth only once and from above within the relevant state-space. The result is important because it validates the theoretical discussion of the previous section.

Furthermore, I produce a comparative statics exercise to examine whether the reduction in churning observed in the data is related to the reduction in growth. I compare steady states of the model under the assumption that the reduction in job reallocation across incumbents is exclusively due to disappearing high growth startups. Evidence of disappearing high growth startups for the U.S. comes from Decker et al. (2016b).

5.1. Calibration

This subsection presents the calibration, which consists in matching selected moments. I pick as many moments as parameters to calibrate, so that I can match them exactly. As data on business dynamism are available from

1978 to 2019, I rely on averages over that period, unless stated otherwise. I divide my discussion in externally calibrated parameters, namely those that have a one-to-one correspondence with a selected moment, and internally calibrated parameters, those that interact with each other to deliver the targeted moment.

Externally Calibrated Parameters

The externally calibrated parameters are summarized in table 1. The employment growth rate, λ, for the U.S. averages 1.6% per year. The parameter ζ is set to 0.55 to match the return to labor in R&D in the U.S. based on NSF data (Mand, 2019). $\beta = 0.98$ is selected to match, anticipating a growth rate of per capita consumption of approximately 2%, a real rate of return of approximately 4% (Gomme et al., 2011).

I further set $\epsilon = 3.9$ to match a markup over marginal cost of 35% (De Ridder et al., 2022).

Hall and Lerner (2010), who review the literature on the returns to R&D, report a widely different ratio of social to private returns to R&D in the various estimations performed over the years. The only consensus seems to be that social returns are substantially larger than private returns. In line with Bloom et al. (2013), I target a ratio of social returns to private returns to knowledge of 2, which requires $\mu = 0.33$. θ is the elasticity of output with respect to knowledge. While Hall and Lerner (2010) reports different estimates from the literature, a value of 0.1 seems like a good compromise between them.

Finally, I estimate from Compustat data the parameters ρ and σ_α. I use all firms with positive sales and R&D values who have at least three consecutive observations (which is the minimum number of consecutive periods needed to pursue the estimation) from 1978 to 2019. I compute their level of technological knowledge using equation (9), and their R&D productivity using equation (12), assuming that the relevant parameters take the values presented in this section. I then estimate the AR(1) process of equation (13), finding $\rho = 0.71$ and $\sigma_\alpha = 1.89$.

Internally Calibrated Parameters

The remaining parameters jointly determine the other targeted moments implied by the model. Some moments are particularly informative of the size of these internally calibrated parameters. For what concerns entrants, the model would ideally require product-level data. Due to data availability

Parameter & target	Symbol & value
Labor force growth	$\lambda = 1.6\%$
Returns to R&D labor	$\zeta = 0.55$
Discount rate	$\beta = 0.98$
Elasticity of substitution between goods	$\epsilon = 3.9$
Persistence R&D productivity	$\rho = 0.71$
Standard deviation R&D productivity shock	$\sigma_\alpha = 1.89$
Returns to knowledge in production	$\theta = 0.1$
Private vs social knowledge	$\mu = 0.33$

Table 1: Externally calibrated parameters.

issues, I rely on establishment-level data. f_E is chosen to match the entry rate of new establishments, which, according to U.S. Census data from the Business Dynamics Statistics database, averages 11.7%. I use the same source to compute an average establishment size of 17 workers, which I use to find the fixed operating cost.

Following Lee and Mukoyama (2015), establishment-level relative productivity of entrants is on average 0.96, from which I determine the average knowledge level of entrants. I instead calibrate its standard deviation to the same standard deviation level as the rest of the market, motivated by evidence from (Sterk et al., 2021) which find that the dispersion of firms does not change substantially throughout their life.

Next, I need to select the parameters χ_α and σ_α^E, that determine the R&D productivity distribution of entrants. I pick the parameter that determines the initial average R&D productivity of entrants to match evidence regarding the contribution of young firms for growth. Specifically, I rely on Foster et al. (2008) estimates of their productivity growth accounting decomposition at a 5-year horizon, and I match the share of productivity growth attributed to entry, which is 24%. Matching evidence from the same type of decomposition presented in other papers, or matching the share of net entry as opposed to entry, would imply a lower average R&D productivity of entrants in my model. I choose the literature's upper bound estimate on the contribution of young firms for growth because I prefer erring in the upward direction than otherwise, given the question I ask in the exercise I present later. Furthermore, the results presented in Foster et al. (2008) include more years of observation and improved techniques for isolating productivity than other highly cited candidates, such as Foster et al. (2001). To identify the standard deviation of the entrants' R&D productivity's underlying normal distribution, instead, I use the exit rate of 1-year-old establishments from the BDS database, which is 26.9%.

Finally, I pick $\overline{\alpha}$ to match a labor productivity growth rate of 1.8%. Importantly, following Bilbiie et al. (2012), I distinguish between data-consistent moments and model-consistent moments when it comes to growth. Note that in this model labor productivity grows because of an increase in average productivity and because of a variety expansion effect. Since the data collection process misses the effect on productivity through variety expansion, the target of 1.8% is for average labor productivity growth.

Parameter	Symbol & value	Main target
Entrant's mean R&D productivity	$\log\chi_\alpha = -6.65$	Entrants' growth contribution: 24%
Entrant's mean initial knowledge	$\log\chi_Z = -2.13$	Average productivity of entrants: 0.96
Standard deviation entrants' R&D productivity	$\sigma_\alpha^E = 4.44$	Entrants' exit rate: 26.9%
St. dev. relative knowledge startups	$\sigma_Z^E = 1.18$	Dispersion employment entrants vs all: 1.0
Fixed operating cost	$\Phi = 3.24$	Average establishment size: 17
Fixed entry cost	$f_E = 3.10$	Establishment entry rate: 11.7%
R&D productivity	$\log\overline{\alpha} = -7.21$	Labor productivity growth: 1.8%

Table 2: Internally calibrated parameters.

To ensure a stationary equilibrium for the firm size distribution, therefore verifying that the calibration delivers the case under which the theoretical results discussed in the previous section apply, I have to verify that firms' growth rates cross the average productivity growth rate only once and from above. Figure 3 illustrates growth rates as a function of relative productivity levels. These results deliver a phase diagram that is qualitatively in line with Figure 2

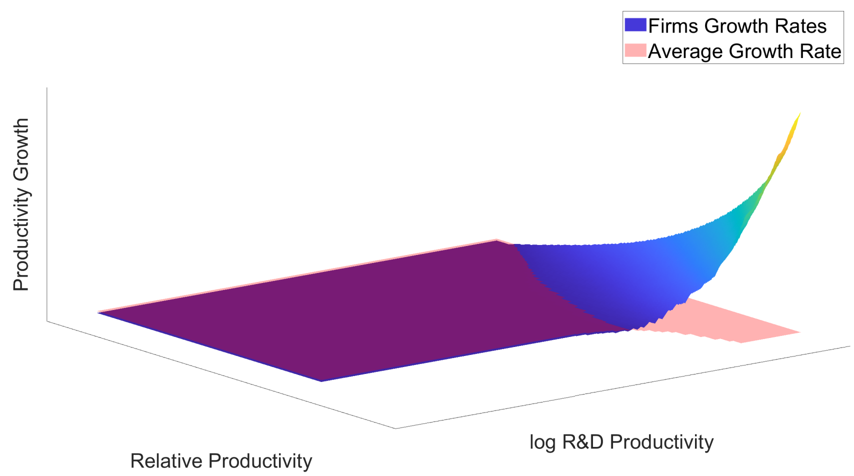

Figure 3: Productivity growth rates and levels.

Note: The graph shows productivity growth rates across the state space. The missing values occur at the combination of R&D productivity and relative productivity at which firms exit the market.

5.2. The Disappearence of Innovative Startups: Consequences for Growth

In this subsection I use the model to answer the following question: can the decline in high-growth startups quantitatively account for the observed coincident changes in turbulence and growth?

Turbulence has been declining in the past few decades in the U.S., with strong discontinuities in job reallocation around the year 2000 (Decker et al., 2016a, 2020). Additionally, Decker et al. (2016b) point out that startups employment growth rates have changed drastically: the right tail of the growth rate distribution has shrunk considerably in that period.

Meanwhile, the growth rate of aggregate productivity has declined considerably from 2005 onwards relative to the previous decade, going back to levels similar to those observed from the mid '70s to the mid '90s (Byrne et al., 2016; Syverson, 2017; Fernald, 2018).

The connection between the two phenomena is puzzling because not all dimensions of turbulence are correlated. The decline in job reallocation, high-growth startups, and productivity growth, occurred around the same time. Instead, declines in firms and establishments entry rates have followed a different trend. The decline in entry started in the early '80s when productivity growth remained unaltered, stopped in the '90s until the mid 2000s while productivity growth increased. Entry then declined permanently during the 2007-2009 recession before stabilizing, after productivity growth had declined according to the studies mentioned above. For this reason, the exercise will focus mostly on dimensions of turbulence other than entry, leaving the puzzle of accounting for the dynamics of entry and growth to future research.

To answer the question, I recalibrate the relevant dimensions of the model until the year 2003. For what concerns the growth rate, I use 2004 as the break year, in line with the model which predicts changes in growth to occur one period after any parameter change, and in line with the literature. This break year maximizes the change in job reallocation across incumbents and growth. I change some targets relative the calibration presented earlier: I target a population growth rate of 1.9%, an entry rate of 12.7%, and a productivity growth rate of 2.1%, which come from the same data until year 2003. Table 3 summarizes the relevant parameter values.

As a relevant non-targeted moment, the model produces a job reallocation rate across incumbents that is approximately 40% of that observed in the data over the same time period. This is a substantial fraction of the overall reallocation rate, considering that differences in productivity across firms are driven by many other different forces, such as the external environment, which are often subject to change (Syverson, 2011).

Parameter	Symbol & value
Entrant's mean R&D productivity	$\log \chi_\alpha = -7.11$
Entrant's mean initial knowledge	$\log \chi_Z = -2.62$
Standard deviation entrants' R&D productivity	$\sigma_\alpha^E = 4.52$
St. dev. relative knowledge startups	$\sigma_Z^E = 1.37$
Fixed operating cost	$\Phi = 3.06$
Fixed entry cost	$f_E = 2.80$
R&D productivity	$\log \overline{\alpha} = -6.60$

Table 3: Internally calibrated parameters until 2003.

The exercise consists in removing high growth startups from the model.

I do that by reducing the parameter σ_α^E, which is the standard deviation of the normal distribution underlying the lognormal draw of startups' R&D productivity. As a result, entrants ability to innovate is lower on average, and exhibits lower dispersion and skewness, which is the simplest and most immediate way to generate what Decker et al. (2016b) document. As the results show that job reallocation across incumbents is quite sensitive to this parameter, I reduce σ_α^E until I match the decline in that moment, which is 2.3 percentage points.

I can then rephrase the question as follows: suppose that all the decline in job reallocation across incumbents is driven by lower σ_α^E, what is the consequence for growth? I do not attempt to provide an explanation for the disappearence of innovative startups, nor I claim that the decline in job reallocation is necessarily driven by that. Instead, I ask whether it is reasonable to expect that the substantial reduction in growth observed could be caused by disappearing high growth startups, regardless of its cause.

The results, presented in table 4 confirm that the connection between lower churning, lower growth, and changes in the composition of entrants is quantitatively credible. After lowering σ_α^E, almost all of the observed decline in growth can be explained by the model. Nevertheless, this parameter change explains very little of the change in entry rates, which, given the discussion proposed earlier in this subsection, may be interpreted as a reasonable description of the data.

Moment	Model	Data
Δ Job reallocation incumbents (target)	-2.3 pp	-2.3 pp
Δ Productivity growth	-0.7 pp	-0.8 pp
Δ Entry	-0.2 pp	-2.2 pp
Δ R&D to GDP	0.01 pp	0.5 pp

Table 4: Results following a disappearence of highly innovative startups.

The reasons for the decline in growth and in job reallocation are quite intuitive. A reduction in the skewness of abilities to innovate ensures that there are fewer good innovators in the market. The presence of good innovators is beneficial for growth due to a few different reasons. First, they directly affect growth because their higher R&D productivity brings up average R&D productivity. As a result, the economy is populated by more good innovators, favoring growth. Second, knowledge spillovers ensure that the more productive a firm is, the faster other firms will grow in absolute terms, holding R&D constant. Third, as good innovators gain market share, they increase the incentives for other firms to perform R&D because of the mechanism illustrated throughout the paper. This last point is part of the explanation for why aggregate R&D does not decline while growth does.

Churning instead declines because a combination of two conditions needs to be true for churning to arise in steady state. First, firms' R&D productivity has to differ. Second, firms have to be away from their no-churning locus (a concept introduced in the previous section). Considering that in the data entrants are generally smaller than incumbents, innovative entrants are on average small, but their R&D productivity implies a much higher size in order for them not to cause any churning. Consequently, these firms contribute directly and disproportionately to churning, as highlighted by the theoretical discussion in the previous section. Furthermore, they also contribute to churning indirectly: as their presence fuels faster average growth, they promote faster shrinkage of firms with a low ability to innovate (which constitutes the majority of firms as R&D productivity is skewed).

Note that the reduction in growth is not accompanied by a similar reduction in R&D spending. The most salient effect of a reduction in σ_α^E is a drop in the average ability to innovate in the market. However, incentives to innovate depend on market share, which by definition aggregates always to 100%. Thus, in the absence of any dramatic change in the distribution of market shares, R&D spending would not be affected considerably: although reducing R&D productivity affects growth in *absolute* productivity per unit of R&D, it may not affect growth in *relative* productivity, which is the only measure of productivity that affects profits.

Although in the data private R&D to GDP increases, this trend is quite smooth throughout the sample period, and does not exhibit discontinuities around the early 2000s when job reallocation and growth change. Therefore, a reasonable interpretation of the results is that the model does not predict any change in R&D following the disappearence of high growth startups, while the time series of R&D seems to agree by not indicating a shift from its sample trend, which is clearly driven by other factors.

6. Summary and Conclusions

This paper presents a unified framework for studying aggregate productivity growth and turbulence. The model features a continuum of monopolistically competitive firms subject to idiosyncratic shocks to their R&D productivity. Because firms perform R&D in-house and competition for market share drives diminishing returns per product in relative terms, R&D investment ultimately declines in relative size, causing endogenous churning. On the other hand, due to increasing returns to scale in absolute terms, the model delivers a positive steady-state productivity growth rate. Therefore, aggregate productivity growth is a turbulent process characterized by movements within the firm size distribution as firms adjust their size optimally by choosing their innovation effort. Entry and exit create life-cycle effects that fuel churning and shape the aggregate growth process by gradually replacing goods sold in the market.

The paper merges and advances two independent streams of literature on firm dynamics (Hopenhayn, 1992) and endogenous growth (Peretto, 1998; Dinopoulos and Thompson, 1998; Young, 1998; Peretto and Connolly, 2007). Specifically, I add endogenous productivity growth to the former, and turbulence to the latter.

In a quantitative analysis, I examine the impact of the disappearance of highly innovative startups, as documented by Decker et al. (2016a) using U.S. data. The results show that churning is highly sensitive to this phenomenon. If the decline in job reallocation among incumbents is solely attributable to this factor, the model can account for nearly the entire slowdown in productivity growth during the same period. However, the disappearance of highly innovative startups is represented in the model through a single parameter adjustment, leaving its underlying causes unexplained. Instead, the model demonstrates that the link between reduced churning and slower productivity growth, driven by a decline in high-growth startups, is plausible.

The framework could prove valuable for several endeavors. For example, it could be employed within the field of firm dynamics that focuses on understanding the effects of friction on resource allocation. The focus would shift from the aggregate productivity level to its growth rate. Additionally, it could prove useful in answering industrial organization questions that have implications for economic growth.

Appendix A. Stationary Model

I present the detrended version of the model, described by the following equations.

First, define $\forall X$, $\widetilde{X}_t = \frac{X_t}{Z_t^\theta}$; $\check{X}_t = \frac{X_t}{N_t^{\frac{1}{\epsilon-1}}}$; $\widehat{X}_t = \frac{X_t}{Z_t^\theta N_t^{\frac{1}{\epsilon-1}}}$.

The production function (9) is:

$$\widetilde{x}_{i,t} = z_{i,t}^\theta l_{x_{i,t}}, \tag{A.1}$$

where the first order condition for production labor is

$$l_{x_{i,t}} = \left[\frac{\epsilon-1}{\epsilon\widehat{w}_t}\left(\frac{\widehat{y}_t}{n_t}\right)^{\frac{1}{\epsilon}} z_{i,t}^{\theta\frac{\epsilon-1}{\epsilon}}\right]^\epsilon, \tag{A.2}$$

where $y_t = \frac{Y_t}{L_t}$and for pricing:

$$\check{p}_{i,t} = \frac{\epsilon}{\epsilon-1}\frac{\widehat{w}_t}{\nu z_{i,t}^\theta}. \tag{A.3}$$

Plugging these into detrended dividend, it be re-expressed as a function of $z_{i,t}$ and $l_{Z_{i,t}}$ only:

$$\widehat{\pi}_{i,t}(z_{i,t}, l_{Z_{i,t}}) = \widehat{w}_t\left[\frac{\epsilon}{(\epsilon-1)} - 1\right]\left[\frac{\epsilon-1}{\epsilon\widehat{w}_t}\left(\frac{\widehat{y}_t}{n_t}\right)^{\frac{1}{\epsilon}} z_{i,t}^{\theta\frac{\epsilon-1}{\epsilon}}\right]^{\frac{\epsilon}{\epsilon-(\epsilon-1)}} - \widehat{w}_t(l_{Z_{i,t}} + \Phi) \tag{A.4}$$

The stationary Bellman equation is:

$$\widehat{V}(z_{i,t}) = \max_{\{l_{Z_{i,t}}\}}\left\{\widehat{\pi}_{i,t}(z_{i,t}, l_{Z_{i,t}}) + (1+g_{t+1}^Z)^\theta(1+\lambda)^{\frac{1}{\epsilon-1}}\left(\frac{n_{t+1}}{n_t}\right)^{\frac{1}{\epsilon-1}}\frac{\widehat{c}_{t+1}}{\widehat{c}_t}\times \frac{1}{1+r_{t+1}}\max\{\mathbb{E}_t\widehat{V}(z_{i,t+1}), 0\}\right\} \tag{A.5}$$

with the knowledge accumulation equation (12), whose stationary version is:

$$z_{i,t} = \frac{z_{i,t-1} + \alpha_{i,t} z_{i,t-1}^\mu l_{Z_{i,t-1}}^\zeta}{(1+g_t^Z)}, \tag{A.6}$$

The entry condition (18) is:

$$\mathbb{E}_t \widehat{v^E}_{i,t}(\alpha_{i,t+1}, z_{i,t+1}) \geq f_E, \tag{A.7}$$

The equilibrium conditions are modified as follows. The labor market clearing (22) becomes:

$$\frac{1}{n_t} = \frac{\int_0^{N_t} (l_{x_{i,t}} + l_{Z_{i,t}}) di}{N_t} + \Phi; \tag{A.8}$$

the law of motion of the number of establishment (19) is now:

$$\frac{n_{t+1}}{n_t} = \frac{1 + \frac{N_{E_t} - N_{X_t}}{N_t}}{1 + \lambda}; \tag{A.9}$$

output (21) is:

$$y_t = c_t + f_E \frac{N_{E_t}}{N_t} n_t; \tag{A.10}$$

Appendix B. Steady State Algorithm

Constuct a grid for the state z (relative knowledge) and the shock α (productivity of R&D) by choosing respectively 180 and 85 grid points. The grid points are spaced in a way to obtain higher concentration for lower values, where non-linearities are present.

Provide an initial guess for the detrended values of wage, output, number of firms and for the growth rate of average knowledge. These are the variables that firms take as given when making their decisions. I use a bisection method to update these guesses. Additionally, I provide an initial guess for the distribution of firms over the firm-specific state variable and shock (relative knowledge and productivity of R&D).

Solve the firm's problem given by the detrended Bellman equation (A.5) via policy function iteration for the R&D labor of firms at each combination of grid points of the two state variables, subject to the constraint (A.6). Furthermore, I compute the value of the firm after dividend payout. If this value is negative, R&D labor is set to 0, as the firm exits the market at the end of the period.

Solve for the expected value of entrants by using the value function computed above. As the value of entrants corresponds to the present value of next period firm value, the firm's decision depends on the expectation of the draw of $\alpha_{i,t+1}$ and $z_{i,t+1}$. This expectation is approximated by a Gauss-Hermitian quadrature with 15 nodes.

At this point, I find the beginning of the period stationary distribution given the guesses for the relevant aggregate variables. This is done by following these steps:

- From the previous period distribution, set the mass of firms at grid points for which firm value is negative to 0. I use the sum of the mass of remaining firms to compute the exit rate, before reweighting the distribution to ensure that the weights of continuing firms sum up to 1.

- Find the new distribution over α, given the old distribution and the law of motion of α. At the same time, find the new distribution of incumbents over $z_{i,t}$. This depends on the old distribution, R&D labor hired in the previous period at given $z_{i,t-1}$ and $\alpha_{i,t-1}$, on $z_{i,t-1}$, on $\alpha_{i,t-1}$.

- Find the distribution of firms that entered in the previous period over the state variable and shock, by drawing $z_{i,t}$ according to equation (16) and $\alpha_{i,t}$ according to equation (17).

- Find the entry rate as the sum of exit rate and population growth rate (the condition required to ensure stationarity in the number of firms, essentially imposing steady state) from equation (A.9).

- Compute the new mass of firms as the weighted average of the mass of incumbents and the mass of entrants, using the entry rate as the weight.

- Iterate until the mass of firms in every grid point is close enough from what it was in the previous iteration.

Finally, the guesses of the aggregate variables need to be updated (I do so by using the bisection method). Find average production and R&D labor using the normalized distribution and the policy functions at each grid point. Compute the values output and number of firms from equations (A.10) and (A.8) respectively. Increase the wage if the left side of equation (A.7) is larger than the right side, and increase the growth rate of average knowledge if the distribution of firms over z is such that the average relative knowledge is larger than 1. Iterate until the values of consumption, number of firms, growth rate of average knowledge, wage, mass of firms over the state and shock and entry rate differ from the values obtained in the previous iteration by less than arbitrary tolerance levels.

The algorithm for a fixed number of firms changes slightly. The wage is decreased if the exit rate is higher than 0, and increased if it is 0.

Appendix C. Proof of Proposition 2

To prove condition (i), I derive the derivative of the optimal R&D investment with respect to relative knowledge. The goal is to show that it is strictly decreasing when $\theta(\epsilon - 1) < 1$. Because the third assumption in the proposition implies that the only relevant time periods are the first two, the derivative obtained through implicit differentiation is:

$$\frac{\partial l_{Z_{i,t}}}{\partial z_{i,t}} = \frac{\frac{1}{1-\zeta} l_{Z_{i,t+1}}^{\frac{\zeta}{1-\zeta}} \frac{B_{t+1}\alpha_{i,t}}{\widehat{w}_t} \zeta \left[\theta\left(\epsilon-1\right)-1\right] \widehat{w}_{t+1}^{1-\epsilon} \left(\frac{\epsilon-1}{\epsilon}\right)^{\epsilon} \frac{\widehat{y}_{t+1}}{n_{t+1}} \theta z_{i,t+1}^{\theta(\epsilon-1)-2}}{\left\{1-\alpha_{i,t} \frac{1}{1-\zeta} l_{Z_{i,t+1}}^{\frac{\zeta}{1-\zeta}} \frac{B_{t+1}\alpha_{i,t}}{\widehat{w}_t} \zeta \left[\theta\left(\epsilon-1\right)-1\right] \widehat{w}_{t+1}^{1-\epsilon} \left(\frac{\epsilon-1}{\epsilon}\right)^{\epsilon} \frac{\widehat{y}_{t+1}}{n_{t+1}} \theta z_{i,t+1}^{\theta(\epsilon-1)-2}\right\}}. \tag{C.1}$$

Notice that, because $\theta(\epsilon - 1) - 1$ is the only factor that can assume negative values, when $\theta(\epsilon - 1) < 1$ the numerator is unambiguously negative, while the denominator is unambiguously positive.

Next, I turn to proving condition (ii). The proof consists of two parts. First, I show that in the limit, as relative knowledge tends to zero, R&D investment tends to infinity. Then, I show that as relative knowledge tends to infinity, R&D investment tends to zero.

Because R&D investment is a function of future relative knowledge, I need to connect current relative knowledge with future relative knowledge. A first useful result is the following:

$$\lim_{z_{i,t}\to\infty} z_{i,t+1} = \frac{z_{i,t} + \alpha_{i,t} l_{Z_{i,t}}^{\zeta}}{1 + g_{t+1}^{Z}} = \infty, \tag{C.2}$$

for any value of $l_{Z_{i,t}}$, given its non-negativity constraint. As $z_{i,t}$ tends to infinity, $z_{i,t+1}$ tends to infinity as well, implying that $l_{Z_{i,t}}$ given in equation

(36) will tend to 0 when $\theta(\epsilon - 1) < 1$. It is helpful here to write equation (12) in growth rates, under the assumption that $\mu = 0$:

$$g_{t+1}^{Z} = \frac{\alpha_{i,t} l_{Z_{i,t}}^{\zeta}}{z_{i,t}}. \tag{C.3}$$

When R&D tends to 0 and relative knowledge tends to infinity, knowledge growth tends to 0, proving the first part of condition (ii).

For the second part of the result, it is sufficient to prove that when relative knowledge tends to 0, R&D does not tend to 0 either, because from equation (C.3) it is immediately visible that when knowledge tends to 0 and R&D is positive, the growth rate of knowledge tends to infinity. I prove this by contraddiction. Suppose that $l_{Z_{i,t}} = 0$, then:

$$\lim_{z_{i,t} \to 0} z_{i,t+1} = \frac{z_{i,t} + \alpha_{i,t} l_{Z_{i,t}}^{\zeta}}{1 + g_{t+1}^{Z}} = 0. \tag{C.4}$$

However, when replacing $z_{i,t+1}$ with 0, equation (36) gives an infinite value of R&D if $\theta(\epsilon - 1) < 1$. Therefore, it is impossible for R&D to be 0 when $z_{i,t}$ tends to 0, implying that the knowledge growth rate must tend to infinity when $z_{i,t}$ tends to 0.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., Kerr, W., 2018. Innovation, Reallocation, and Growth. American Economic Review 108, 3450–3491. doi:10.1257/aer.20130470.

Acemoglu, D., Cao, D., 2015. Innovation by entrants and incumbents. Journal of Economic Theory 157, 255–294. doi:10.1016/j.jet.2015.01.001.

Acemoglu, D., Ventura, J., 2002. The World Income Distribution. The Quarterly Journal of Economics 117, 659–694. arXiv:2696438.

Aghion, P., Howitt, P., 1992. A Model of Growth Through Creative Destruction. Econometrica 60, 323–351. doi:10.2307/2951599, arXiv:2951599.

Akcigit, U., Ates, S.T., 2019. What Happened to U.S. Business Dynamism? Working Paper 25756. National Bureau of Economic Research. doi:10.3386/w25756.

Akcigit, U., Kerr, W.R., 2018. Growth through Heterogeneous Innovations. Journal of Political Economy 126, 1374–1443. doi:10.1086/697901.

Audretsch, D., Klomp, L., Santarelli, E., Thurik, A., 2004. Gibrat's Law: Are the Services Different? Review of Industrial Organization 24, 301–324. doi:10.1023/B:REIO.0000038273.50622.ec.

Bilbiie, F.O., Ghironi, F., Melitz, M.J., 2012. Endogenous Entry, Product Variety, and Business Cycles. Journal of Political Economy 120, 304–345. doi:10.1086/665825.

Bloom, N., Schankerman, M., Reenen, J.V., 2013. Identifying Technology Spillovers and Product Market Rivalry. Econometrica 81, 1347–1393. doi:10.3982/ECTA9466.

Bottazzi, G., Cefis, E., Dosi, G., Secchi, A., 2007. Invariances and Diversities in the Patterns of Industrial Evolution: Some Evidence from Italian Manufacturing Industries. Small Business Economics 29, 137–159. doi:10.1007/s11187-006-0014-y.

Brown, C., Haltiwanger, J., Lane, J., 2008. Economic Turbulence: Is a Volatile Economy Good for America? University of Chicago Press.

Byrne, D.M., Fernald, J.G., Reinsdorf, M.B., 2016. Does the United States Have a Productivity Slowdown or a Measurement Problem? Brookings Papers on Economic Activity 2016, 109–182. doi:10.1353/eca.2016.0014.

Cao, D., Hyatt, H.R., Mukoyama, T., Sager, E., 2017. Firm Growth through New Establishments .

De Ridder, M., 2024. Market Power and Innovation in the Intangible Economy. American Economic Review 114, 199–251. doi:10.1257/aer.20201079.

De Ridder, M., Grassi, B., Morzenti, G., 2022. The Hitchhiker's Guide to Markup Estimation .

Decker, R.A., Haltiwanger, J., Jarmin, R.S., Miranda, J., 2016a. Declining Business Dynamism: What We Know and the Way Forward. American Economic Review 106, 203–207. doi:10.1257/aer.p20161050.

Decker, R.A., Haltiwanger, J., Jarmin, R.S., Miranda, J., 2016b. Where has all the skewness gone? The decline in high-growth (young) firms in the U.S. European Economic Review 86, 4–23. doi:10.1016/j.euroecorev.2015.12.013.

Decker, R.A., Haltiwanger, J., Jarmin, R.S., Miranda, J., 2020. Changing Business Dynamism and Productivity: Shocks versus Responsiveness. American Economic Review 110, 3952–3990. doi:10.1257/aer.20190680.

Dinopoulos, E., Thompson, P., 1998. Schumpeterian Growth Without Scale Effects. Journal of Economic Growth 3, 313–335. doi:10.1023/A:1009711822294.

Dosi, G., 1988. Sources, Procedures, and Microeconomic Effects of Innovation. Journal of Economic Literature 26, 1120–1171. arXiv:2726526.

Ericson, R., Pakes, A., 1995. Markov-Perfect Industry Dynamics: A Framework for Empirical Work. The Review of Economic Studies 62, 53–82. doi:10.2307/2297841, arXiv:2297841.

Fernald, J., 2018. Is Slow Productivity and Output Growth in Advanced Economies the New Normal? International Productivity Monitor , 138–148.

Foster, L., Haltiwanger, J., Krizan, C., 2001. Aggregate Productivity Growth: Lessons from Microeconomic Evidence, in: New Developments in Productivity Analysis. University of Chicago Press, Chicago. Edited by C.R. Hulten, E.R. Dean, and M.J. Harper, pp. 303–363. doi:10.3386/w6803.

Foster, L., Haltiwanger, J., Syverson, C., 2008. Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? American Economic Review 98, 394–425. doi:10.1257/aer.98.1.394.

Garcia-Macia, D., Hsieh, C.T., Klenow, P.J., 2019. How Destructive Is Innovation? Econometrica 87, 1507–1541. doi:10.3982/ECTA14930.

Gomme, P., Ravikumar, B., Rupert, P., 2011. The return to capital and the business cycle. Review of Economic Dynamics 14, 262–278. doi:10.1016/j.red.2010.11.004.

Ha, J., Howitt, P., 2007. Accounting for Trends in Productivity and R&D: A Schumpeterian Critique of Semi-Endogenous Growth Theory. Journal of Money, Credit and Banking 39, 733–774. doi:10.1111/j.1538-4616.2007.00045.x.

Hall, B.H., Lerner, J., 2010. Chapter 14 - The Financing of R&D and Innovation, in: Hall, B.H., Rosenberg, N. (Eds.), Handbook of the Economics of Innovation. North-Holland. volume 1 of *Handbook of The Economics of Innovation, Vol. 1*, pp. 609–639. doi:10.1016/S0169-7218(10)01014-2.

Hopenhayn, H.A., 1992. Entry, Exit, and firm Dynamics in Long Run Equilibrium. Econometrica 60, 1127–1150. doi:10.2307/2951541, `arXiv:2951541`.

Hopenhayn, H.A., 2014. Firms, Misallocation, and Aggregate Productivity: A Review. Annual Review of Economics 6, 735–770. doi:10.1146/annurev-economics-082912-110223.

Klette, T.J., Kortum, S., 2004. Innovating Firms and Aggregate Innovation. Journal of Political Economy 112, 986–1018. doi:10.1086/422563.

Laincz, C.A., 2009. R&D subsidies in a model of growth with dynamic market structure. Journal of Evolutionary Economics 19, 643–673. doi:10.1007/s00191-008-0114-8.

Laincz, C.A., Peretto, P.F., 2006. Scale effects in endogenous growth theory: An error of aggregation not specification. Journal of Economic Growth 11, 263–288. doi:10.1007/s10887-006-9004-9.

Lee, Y., Mukoyama, T., 2015. Entry and exit of manufacturing plants over the business cycle. European Economic Review 77, 20–27. doi:10.1016/j.euroecorev.2015.03.011.

Lentz, R., Mortensen, D.T., 2008. An Empirical Model of Growth Through Product Innovation. Econometrica 76, 1317–1373. doi:10.3982/ECTA5997.

Lentz, R., Mortensen, D.T., 2016. Optimal growth through product innovation. Review of Economic Dynamics 19, 4–19. doi:10.1016/j.red.2015.12.002.

Luttmer, E.G.J., 2007. Selection, Growth, and the Size Distribution of Firms. The Quarterly Journal of Economics 122, 1103–1144. doi:10.1162/qjec.122.3.1103.

Madsen, J.B., 2008. Semi-endogenous versus Schumpeterian growth models: Testing the knowledge production function using international data. Journal of Economic Growth 13, 1–26. doi:10.1007/s10887-007-9024-0.

Mand, M., 2019. On the cyclicality of R&D activities. Journal of Macroeconomics 59, 38–58. doi:10.1016/j.jmacro.2018.10.008.

Melitz, M.J., Polanec, S., 2015. Dynamic Olley-Pakes productivity decomposition with entry and exit. The RAND Journal of Economics 46, 362–375. doi:10.1111/1756-2171.12088.

Naudé, W., 2022. From the entrepreneurial to the ossified economy. Cambridge Journal of Economics 46, 105–131.

Peretto, P.F., 1998. Technological Change, Market Rivalry, and the Evolution of the Capitalist Engine of Growth. Journal of Economic Growth 3, 53–80. doi:10.1023/A:1009722031825.

Peretto, P.F., 1999. Cost reduction, entry, and the interdependence of market structure and economic growth. Journal of Monetary Economics 43, 173–195. doi:10.1016/S0304-3932(98)00040-3.

Peretto, P.F., 2018. Robust endogenous growth. European Economic Review 108, 49–77. doi:10.1016/j.euroecorev.2018.06.007.

Peretto, P.F., Connolly, M., 2007. The Manhattan Metaphor. Journal of Economic Growth 12, 329–350. doi:10.1007/s10887-007-9023-1.

Peretto, P.F., Smulders, S., 2002. Technological distance, growth and scale effects. The Economic Journal 112, 603–624. doi:10.1111/1468-0297.00732.

Peters, M., 2020. Heterogeneous Markups, Growth, and Endogenous Misallocation. Econometrica 88, 2037–2073. doi:10.3982/ECTA15565.

Restuccia, D., Rogerson, R., 2017. The Causes and Costs of Misallocation. Journal of Economic Perspectives 31, 151–174. doi:10.1257/jep.31.3.151.

Sterk, V., Sedláček, P., Pugsley, B., 2021. The Nature of Firm Growth. American Economic Review 111, 547–579. doi:10.1257/aer.20190748.

Sutton, J., 1997. Gibrat's Legacy. Journal of Economic Literature 35, 40–59. arXiv:2729692.

Syverson, C., 2011. What Determines Productivity? Journal of Economic Literature 49, 326–365. doi:10.1257/jel.49.2.326.

Syverson, C., 2017. Challenges to Mismeasurement Explanations for the US Productivity Slowdown. Journal of Economic Perspectives 31, 165–186. doi:10.1257/jep.31.2.165.

Thompson, P., 2001. The Microeconomics of an R&D-Based Model of Endogenous Growth. Journal of Economic Growth 6, 263–283. doi:10.1023/A:1012761811439.

Young, A., 1998. Growth without Scale Effects. Journal of Political Economy 106, 41–63. doi:10.1086/250002.