Turbulent Growth: Business Dynamism and Aggregate Productivity*

Filippo Massari†
Department of Economics
Fairfield University

November 8, 2023

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 market share through innovation.

When firms increase their technological knowledge, they gain market share by lowering their relative price, thus reducing the marginal value of further gains in market share. This leads to the emergence of diminishing returns in relative terms. Therefore, incentives to innovate decline in relative size, all else constant, generating churning endogenously as mean-reversion. This mechanism delivers a stationary, non-degenerate, and endogenous firm-size distribution dependent on R&D. Meanwhile, constant returns to the cumulative factor (technological knowledge) drive a trendless aggregate growth rate determined by R&D. Endogenous entry and exit entail selection effects that shape the characteristics of the firm population, and generate a firm life cycle, affecting R&D, thus growth.

(*JEL*: E19, L11, O14, O31, O41)

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

*Acknowledgement: I am grateful to my advisors Pietro Peretto and Giuseppe Fiori for their constant guidance and encouragement. I have also benefited greatly from various discussions with Domenico Ferraro, James Nason, Daisoon Kim, Soroush Ghazi, and Ayse Kabukcuoglu Dur. I thank all parcipants to the Macroeconomics Lunch Seminars at North Carolina State University and seminar participants at Fairfield University, the University of Milan, and Drexel University. Lastly, I acknowledge the technical support of the High-Performance Computing Services at North Carolina State University. All remaining errors are my own.

†Dolan School of Business, Fairfield University, 1073 North Benson Road, Fairfield, CT 06824. E-mail: fmassari@fairfield.edu. Website: http://filippomassari.com/.

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 formation, while lack of it may ultimately lead to firm death. Intuitively, aggregate productivity growth is a *turbulent process* intrinsically connected to business dynamism.

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."

Why should growth theorists be concerned with turbulence? In this paper, I study aggregate productivity growth starting from the processes of firm formation, firm growth, and firm death that drive key aspects of turbulence. The goal is to take seriously widely documented microeconomic phenomena, such as a negative correlation of firms' employment growth with initial relative size (Sutton, 1997; Caves, 1998; Audretsch et al., 2004), and describe their interdependence with macroeconomic variables of interest.

The motivation for this endeavor is two-fold. First, the need for an integrated framework where the various dimensions of turbulence interact with the growth process is evident in the field of firm dynamics. When Restuccia and Rogerson (2017, p. 168) identify 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." This paper delivers what they ask. As a result, future research can concentrate on understanding the sources of differences in income growth patterns across countries and time periods, in addition to the usual focus on income levels.

Second, the joint study of aggregate productivity growth and turbulence poses a non-trivial theoretical challenge. Delivering sustained economic growth endogenously requires constant or increasing returns to scale. However, when allowing for R&D differentials across firms, increasing 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.

Hence, two relevant questions arise: What prevents larger firms from monopolizing the market over time? Is endogenous growth theory robust to introducing forces that drive heterogeneity in firms' investment decisions, thus endogenizing turbulence in all dimensions?

The framework I propose to answer these questions builds on the endogenous growth and the firm dynamics literature. A continuum of firms,[1] produces unique but substitutable goods (monopolistic competition). They perform innovation in-house to enhance their stock of technological knowledge through which they reduce their marginal cost of production, hence their price. By

[1] 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.

reducing their price, firms steal market share from their competitors. Entry and exit decisions are endogenous. In this framework, the fundamental source of heterogeneity is an idiosyncratic stochastic R&D productivity (ability to innovate) drawn from a common distribution. Aggregate variables evolve deterministically.

This setup brings about three main implications that contribute to our understanding of productivity growth.

First, churning is at least partly endogenous and a fundamental property of the aggregate growth process. In particular, I uncover a mechanism through which R&D differs across firms generating differences in growth rates, thus causing churning. The possibility of stealing market share from others provides incentives to innovate. Simultaneously, this competitive force gives rise to *diminishing economic returns*, as 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 market share giving rise to negative scale dependence — the negative correlation between firms' growth rates and their relative size — in the form of mean-reversion.

Is it, therefore, the case that the aggregate productivity growth rate will tend to 0 eventually? No. As these diminishing returns are in *relative* terms, it is possible to accommodate increasing returns to scale in *absolute* terms. Hence, firm heterogeneity in R&D investment does not necessarily disrupt sustained growth since constant returns to the growth-driving factor can be preserved. As a result, the aggregate productivity growth rate depends on firms' optimal R&D decisions as in other fully endogenous growth models.

The theoretical ingredients responsible for this scale dependence are standard: (i) in-house R&D; (ii) an idiosyncratic shock to the innovation process that is the root cause of heterogeneity in firms' knowledge stock and market share; and (iii) imperfect substitutability between goods, which introduces demand-driven 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 relative price. This exerts downward pressure on the return to further market share expansion, thus reducing the incentive to innovate and grow even larger.

A second related contribution concerns the firm-size distribution. Through the economic mechanism illustrated above, I point to a novel stationarity source. As firms' return to R&D investment, all else constant, declines in their level of relative knowledge, this model gives rise to mean-reversion in relative knowledge (hence in size). In this way, it introduces a force that prevents firms from perpetually growing bigger or shrinking. My model produces a firm-size distribution arising from firms' dynamic optimization that is non-degenerate and has a bounded support determined endogenously. 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.

Third, the model characterizes aggregate productivity growth as a process of gradual product replacement through endogenous entry and exit. Firms are subject to an idiosyncratic persistent shock to their R&D productivity, namely their ability to innovate. Entrants hit by a good initial draw and a series of positive shocks thereafter innovate and grow in size. Eventually, their R&D investment will become just enough to allow them to maintain that size. However, as shocks are drawn from a common distribution, eventually, they will be hit by a sequence of average and negative shocks, meaning that the quality of their ideas will revert to the mean. Absent any new good shock, they will start losing market share gradually at the expense of better innovators until it is no longer profitable to remain in the market.

Arguably, the most significant contribution is the framework, which allows studying entry, exit, churning, the number of goods, the firm-size distribution, and aggregate productivity growth jointly. This paper merges the literature on firm dynamics with the endogenous growth literature, specifically Hopenhayn (1992) and Peretto and Connolly (2007).

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 the aforementioned Restuccia and Rogerson (2017), which has recently devoted much attention to resource allocation and the aggregate productivity level.

As in Peretto and Connolly (2007), the theory 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. While that model focuses on the symmetric equilibrium, my model introduces an idiosyncratic shock to firms' R&D outcome, thus delivering a non-degenerate firm-size distribution and richer firm dynamics by giving rise to churning and allowing for simultaneous entry and exit.

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). Notable examples of general equilibrium analyses that give rise to endogenous aggregate productivity growth date back to Thompson (2001) and Laincz (2009).[2] 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. 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.

Finally, delivering a stationary and non-degenerate firm-size distribution has so far been a tough theoretical challenge. My model includes an endogenous mechanism that ensures that firms endowed with a higher ability to innovate do not monopolize the market over time. Other models adopt working assumptions aimed at preserving the distribution's stationarity so that they can study simultaneously aggregate productivity growth and firm dynamics without mathemat-

[2]The relevance of endogenous growth models with heterogeneous producers is corroborated by the proliferation of studies, primarily based on Klette and Kortum (2004), that focus on understanding a variety of phenomena or informing policy (Acemoglu et al., 2018; Mukoyama and Osotimehin, 2019; Akcigit and Kerr, 2018). These papers motivate heterogeneity by pointing out that entrants and mature firms have different abilities to innovate. As a result, heterogeneity plays a role through a selection effect only: parameter changes that affect the entry and exit rates modify the average ability to innovate in the economy, thus R&D incentives, and growth. My model includes this aspect, but it goes a step further as different firm-size distributions imply different incentives to innovate.

ical complications. They either rely on entrants replacing firms that would otherwise become too small, or an exogenous death shock that prevents the most successful firm from becoming too large and monopolizing the market (Klette and Kortum, 2004; Acemoglu and Cao, 2015; Cao et al., 2017; Acemoglu et al., 2018; Akcigit and Ates, 2019). The reason why my model includes both entry and exit is that they are phenomena of interest. However, they are always endogenous outcomes of firms' decisions, and their presence is unnecessary to maintain the firm-size distribution stationary. Other models use knowledge spillovers that facilitate the growth of smaller firms (Thompson, 2001; Laincz, 2009), an aspect included in my model as well.

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 when entry and exit are endogenous. Section 5 provides a quantitative estimation of the relevance of entry, exit, and churning for growth. 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.[3] 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.

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.[4] Finally, at the end of the period, it decides to exit if its continuation value is negative.

Static Problem: Dividends

[3] 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.

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

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 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:

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

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:

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

$0 < \mu < 1$ 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 K_t 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. It takes up the following form:

$$K_t = \frac{1}{N_t} \int_0^{N_t} Z_{i,t} di \equiv Z_t. \tag{13}$$

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.[5]

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, 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:

$$\log \alpha_{i,t} = (1-\rho) \log \overline{\alpha} + \rho \log \alpha_{i,t-1} + \xi_{i,t}, \qquad \begin{array}{c} \xi_{i,t} \sim N(0, \sigma_{\xi}) \\ 0 \leq \rho < 1 \end{array} \tag{14}$$

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:

$$\max_{\{l_{Z_{i,t+h}}, Z_{i,t+1+h}\}_{h=0}^{\infty}} V_{i,t} = \pi_{i,t}(l_{Z_{i,t}}, Z_{i,t}) + \max \left\{ \mathbb{E}_t \sum_{h=1}^{\infty} \prod_{q=1}^{h} \frac{1}{1+r_{t+q}} \pi_{i,t+h}(l_{Z_{i,t+1}}, Z_{i,t+1}), 0 \right\} \tag{15}$$

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

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

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

$$V(Z_{i,t},\alpha_{i,t}) = \max_{\{l_{Z_{i,t}}\}} \left\{ \pi_{i,t}(Z_{i,t}, l_{Z_{i,t}}) + \frac{1}{1+r_{t+1}} \max\{\mathbb{E}_t V(Z_{i,t+1},\alpha_{i,t+1}),0\} \right\} \tag{16}$$

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:

$$Z^d_{t+1} \sim Lognormal(\chi_Z, \sigma^E_Z) Z_{t+1}. \tag{17}$$

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

$$\alpha^d_{t+1} \sim Lognormal(\chi_\alpha, \sigma^E_\alpha). \tag{18}$$

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

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:

$$v^E_t = \mathbb{E}_t V(\alpha^d_{t+1}, Z^d_{t+1}) \geq Z^\theta_t N_t^{\frac{1}{\epsilon-1}} f_E, \tag{19}$$

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 (19);
- a law of motion of firms:
$$N_{t+1} = N_t + N_{E_t} - N_{X_t}; \tag{20}$$
- 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:

$$s_t = \int_0^{N_t - N_{X_t}} v_{i,t} di + \int_0^{N_{E_t}} v_{i,t}^E di. \tag{21}$$

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

$$Y_t = c_t L_t + Z_t^{\theta} N_t^{\frac{1}{\epsilon-1}} f_E N_{E_t}. \tag{22}$$

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

$$L_t = \int_0^{N_t} (l_{x_{i,t}} + l_{Z_{i,t}}) di + \Phi N_t. \tag{23}$$

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

$$z_{i,t} = \frac{Z_{i,t}}{Z_t}, \tag{24}$$

and the number of firms per capita

$$n_t = \frac{N_t}{L_t}, \tag{25}$$

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:

$$A_{i,t} = Z_{i,t}^{\theta}. \tag{26}$$

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:

$$\frac{Y_t}{L_t} = \underbrace{\underbrace{N_t^{\frac{1}{\epsilon-1}}}_{\substack{\text{Tech.}\\ \text{breadth}}} \quad \underbrace{A_t}_{\substack{\text{Tech.}\\ \text{depth}}} \quad \underbrace{\left\{ \frac{1}{N_t} \int_0^{N_t} \left(S_{i,t} N_t \, a_{i,t} \right)^{\frac{\epsilon-1}{\epsilon}} di \right\}^{\frac{\epsilon}{\epsilon-1}}}_{\substack{\text{Allocative}\\ \text{efficiency}}}}_{\text{Aggregate productivity}} \quad \underbrace{\frac{L_{x_t}}{L_t}}_{\text{Production effort}} , \tag{27}$$

where $S_{i,t}$ is the production labor share for firm i. 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 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:

$$A_t = \int_0^{N_t} S_{i,t} A_{i,t} di, \tag{28}$$

As I will clarify later, this definition of average productivity is useful to understand the role that heterogeneity plays in shaping the aggregate productivity level. Coincidentally, this definition is the one 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:

$$a_{i,t} = \frac{A_{i,t}}{A_t}. \tag{29}$$

This element describes the allocative efficiency of the economy. The distribution of individual productivities matters for aggregate productivity for two reasons, both having to do with aggregation. First, 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 (28) can be re-expressed as:

$$\int_0^{N_t} S_{i,t} a_{i,t} di = 1. \tag{30}$$

The term labeled *allocative efficiency* above would therefore equal 1 under a symmetric equilibrium, or in a model with additive aggregation of goods *and* a production function linear in labor. It follows that the term in bracket in equation (27) 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 (27):

$$1 + g^{productivity} = \underbrace{\left[\frac{n_t(z_t, \alpha_t)}{n_{t-1}(z_{t-1}, \alpha_{t-1})}(1+\lambda)\right]^{\frac{1}{\epsilon-1}}}_{\text{semi-endogenous}} \underbrace{\left(1 + g_t^A(z_{t-1}, \alpha_{t-1})\right)}_{\text{average productivity}} \underbrace{\left\{\frac{\frac{1}{N_t}\int_0^{N_t} \left(S_{i,t} N_t \, a_{i,t}\right)^{\frac{\epsilon-1}{\epsilon}} di}{\frac{1}{N_{t-1}}\int_0^{N_{t-1}} \left(S_{i,t-1} N_{t-1} \, a_{i,t-1}\right)^{\frac{\epsilon-1}{\epsilon}} di}\right\}^{\frac{\epsilon}{\epsilon-1}}}_{\text{change in the distribution}}, \tag{31}$$

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 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).

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

3.1 Growth Rate Differentials

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 stationarity and non-degeneracy of the distribution.

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." Similarly, Acemoglu and Cao (2015) claim that the model presented requires a specific type of imitative entry to prevent firms from falling too far behind and preserving the distribution's stationarity. A model that allows both churning within the distribution alone and entry and exit is required if the goal is to understand how heterogeneity and the different dimensions of turbulence interact with the aggregate growth process.

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). I fix the steady state number of firms and remove entry and exit 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.

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

$$g_{t+1}^{A_i} \approx \alpha_{i,t+1} z_{i,t}^{\mu-1} l_{Z_{i,t}} (z_{i,t}, \mathbb{E}_t \alpha_{i,t+1})^{\zeta}. \tag{32}$$

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 of it.

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:

$$l_{Z_{i,t}}^{1-\zeta} = \frac{1}{1+r_{t+1}} \frac{\alpha_{i,t} z_{i,t}^{\mu}}{w_t \left(1+g_{t+1}^{Z}\right)} \mathbb{E}_t \left[\frac{w_{t+1} l_{Z_{i,t+1}}^{1-\zeta}}{\zeta \alpha_{i,t+1} z_{i,t+1}^{\mu}} + \frac{\delta \pi_{i,t+1}}{\delta z_{i,t+1}} + \frac{\mu w_{t+1} l_{Z_{i,t+1}}}{\zeta z_{i,t+1}} \right]. \tag{33}$$

The equation 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 effectiveness 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, η, 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.

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 (1990), as ϵ tends toward infinity. In theat 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? For example, differences in estimates of the Gibrat coefficients found 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.

The calibration illustrated below delivers the parametrizations needed for stationarity: firms' productivity growth rates are decreasing in their relative level, conditional on R&D productivity. The implication is that for any R&D productivity level, relative productivity growth as a function of its level intersects 0 only once and from above. Figure 1 illustrates this relationship for selected levels of R&D productivity.

Figure 2 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.

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

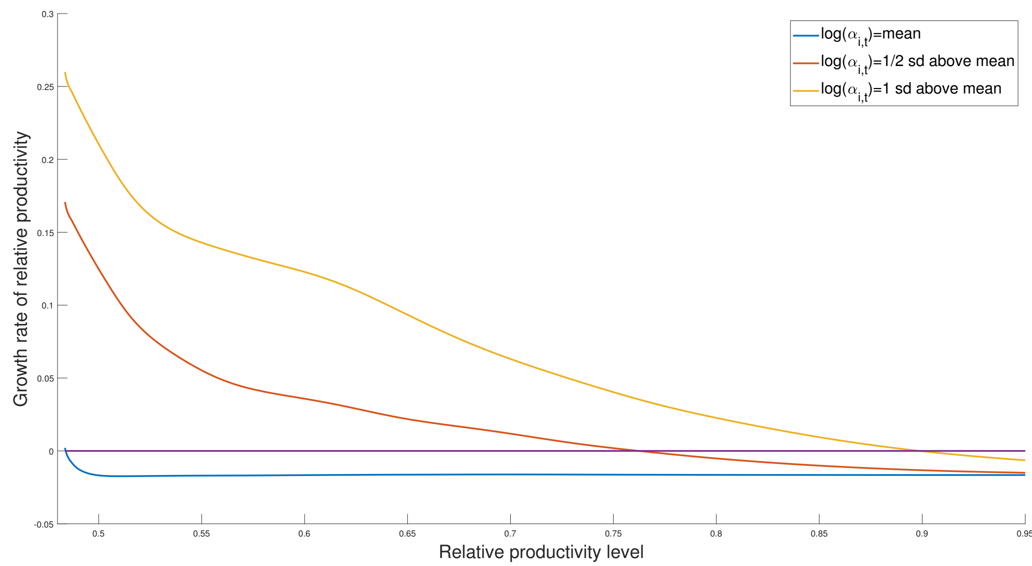

Figure 1: Productivity growth rates and levels.

Note: Relative productivity growth rates and levels for mean R&D productivity, half a standard deviation, and one standard deviation above.

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, a condition that is verified as mentioned above.

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. 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 2 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, R&D investment is strictly increasing in R&D productivity, implying that more innovative firms grow

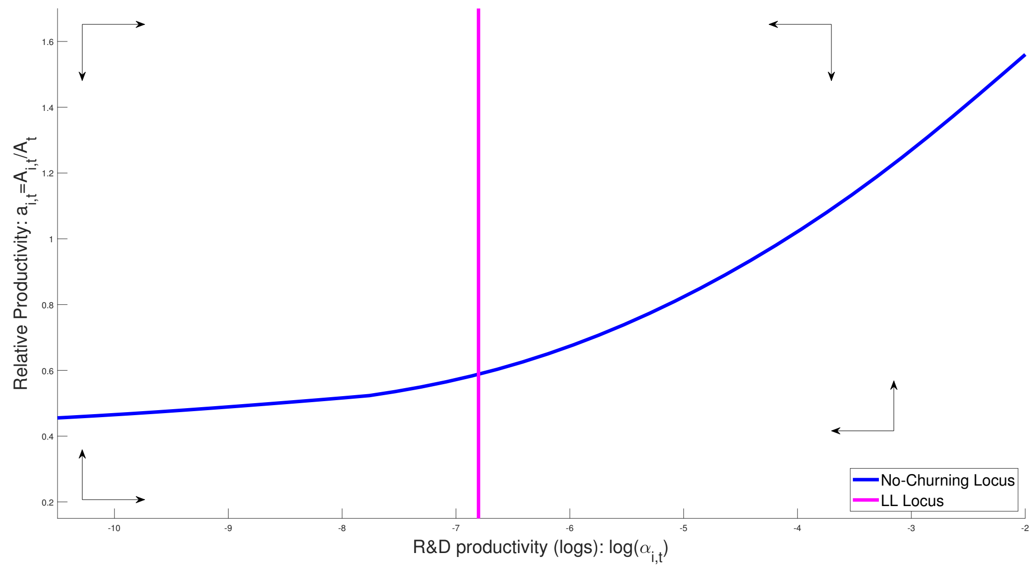

Figure 2: Turbulence, firm growth, and stability.

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.

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.

Finally, Figure 3 plots the counter-cumulative distribution function (in logs) observed in the simulated model compared to a Pareto distribution with a coefficient estimated from the simulated model. The distribution generated by the model is a downward sloping curve for mid-sized firms, where differences in R&D productivity are translated into differences in size. However, it deviates significantly from Pareto for very small and very big firms. As Rossi-Hansberg and Wright (2007) point out, this outcome directly results from the negative dependence of firms' growth rates on their relative size. As firms reach a large relative size, the forces of attraction

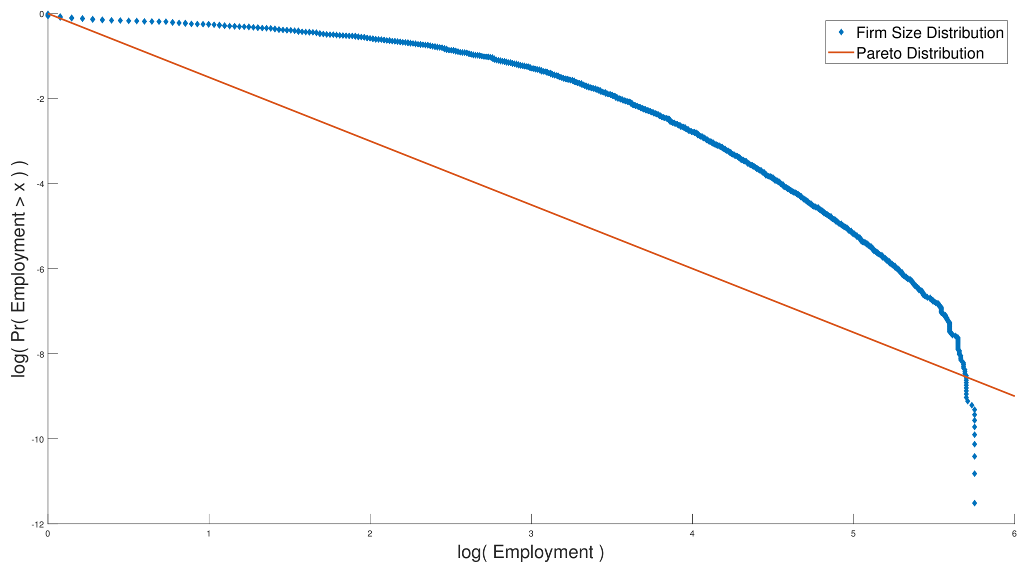

Figure 3: Firm size fdstribution.

Note: Comparison between a counter cumulative distribution function (in logarithmic scale) for the firm sizes generated by the model and for a fitted Pareto distribution.

are too strong for them to grow further, thus leading the right tail of the distribution to collapse quickly. Similarly, when firms become very small relative to others, the knowledge spillover and incentive to innovate are high enough to prevent them from shrinking any further. Notably, this shape is qualitatively in line with the one shown in Figure 5 of Rossi-Hansberg and Wright (2007) for manufacturing establishments in the US.

4 Entry and Exit

In this section, I relax some of 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. 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 are strictly decreasing in relative productivity, conditional on R&D productivity — 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 presents the implications of adding entry and exit.

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 the model's predicted life cycle.

Figure 4 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.

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 4 shows that all firms expect

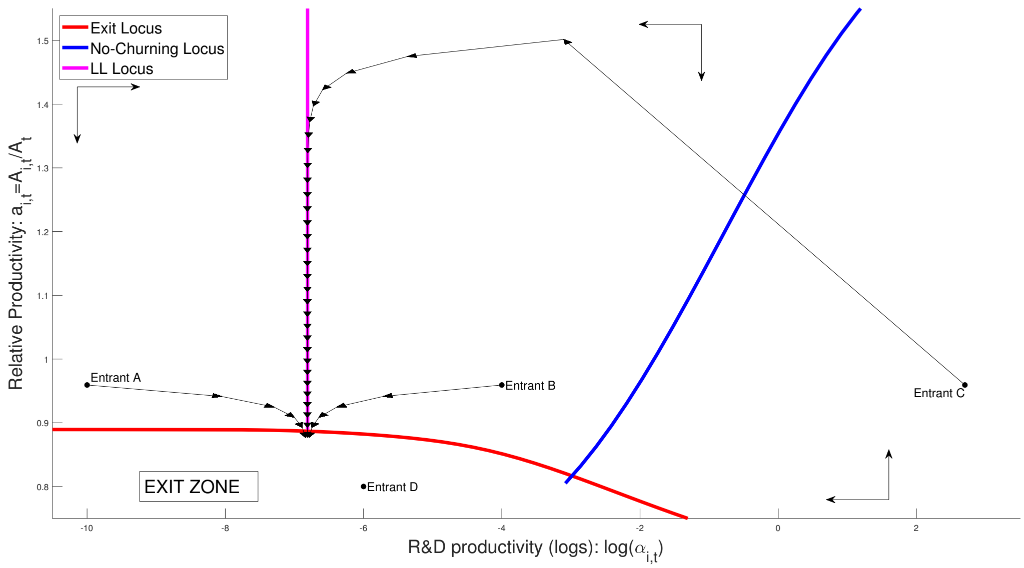

Figure 4: The firm expected life cycle.

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 three startups that differ in their initial R&D productivity draw.

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 average relative productivity level for entrants.

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 US, which exits the market in less than a decade (based on US 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 US Census data, these firms are relatively common as about 20% of the population exits the market within their first year of operation.

5 Turbulence, the Firm Size Distribution, and Their Quantitative Relevance for Growth

How much of the growth process is lost when focusing only on understanding average firm growth? This section aims to answer this question. In particular, I will compare the economy's growth rate delivered by the calibrated model with two alternatives. First, I will look at the growth rate of the model when entry, and consequently exit, is excluded. Then, I will compute the growth rate when the idiosyncratic shocks are shut down, causing the absence of churning.The primary purpose of this exercise is to shed light on the most relevant effects that characterize the interaction between growth, turbulence, and the firm size distribution.

The first subsection illustrates the calibration, while the second shows the quantitative results.

5.1 Calibration

This subsection presents the calibration, which consists in matching selected moments. 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

These externally calibrated parameters are summarized in table 3. The labor force growth rate, λ, for the US is 1% per year. The parameter η is set to 0.55 to match the return to labor in R&D in the US based on NSF data (Mand, 2019). Following Lee and Mukoyama (2015), establishment-level relative productivity of entrants is on average 0.96. $\beta = 0.987$ is selected to match, anticipating a growth rate of per capita consumption of 2.6%, 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.

Table 1: Externally calibrated parameters.

Parameter	Symbol	Target
Labor force growth	λ	1%
Returns to R&D labor	η	0.55
Discount rate	β	real interest rate: 4%
Elasticity of substitution between goods	ϵ	35% markup
Relative knowledge of entrants	χ_z	relative productivity: 0.96
Returns to knowledge in production	θ	0.1
Private vs social knowledge	μ	Social to private returns to knowledge: 2

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. The steady state entry and exit rates depend largely on the productivity of R&D of incumbents relative to entrants. χ_α almost uniquely determines the average productivity growth of entrants. Huergo and Jaumandreu (2004) estimate that in Spanish manufacturing firms in selected industries in the '90s, entrants grow on average faster than the entire population of firms by 2 percentage points. f_E is chosen to match the entry rate of new establishments, which, according to US Census data, averaged 11.6%. Moreover, the parameter σ_α^E, which regulates the dispersion of the initial draw of the average productivity of R&D of entrants, is picked to match the average exit rate, 26.9%, of establishments younger than one year in the US in the same period. Finally, the standard deviation of relative knowledge for entrants is calibrated to match the ratio between the standard deviation of employment for startups and their standard deviation the following year. In the US economy, this value is around 1 (Sterk et al., 2021).

Equation (23) shows that the parameter Φ affects the average size of establishments, which averages approximately 17, implying $n_t = 0.059$ in steady state.

The parameters ρ and σ_α in addition to $\overline{\alpha}$ affect the average productivity of R&D of incumbents in the economy. I target a per capita consumption growth rate of 2.6%. At the same time, the combination of these three parameters regulates the the persistence of firm growth and of firm employment levels. Firm sales' growth rates have an AR(1) coefficient of 0.22 (Dosi et al., 2019). The AR(1) coefficient for log employment is 0.97 (Sterk et al., 2021).

5.2 The Quantitative Results

This subsection illustrates the quantitative relevance of the dimensions of turbulence discussed in the paper.

To isolate each effect, first, I consider a model without net entry (defined as "churning and entry=exit" in the table), achieved by eliminating population growth, $\lambda = 0$. Then, I shut down

Table 2: Internally calibrated parameters.

Parameter	Symbol	Value
Entrant's mean R&D productivity	$\log \chi_\alpha$	-10.2
Standard deviation entrants' R&D productivity	σ_α^E	5.8
Fixed operating cost	Φ	3.0
Fixed entry cost	f_E	7.2
Shock persistence	ρ	0.39
R&D productivity	$\log \bar{\alpha}$	-6.8
St. dev. relative knowledge startups	σ_Z^E	1.3
St. dev. draw of R&D productivity shock	σ_ξ	2.4

Table 3: Targets for internally calibrated parameters.

Parameter	Source	Value
Average plant size	BDS	17
Entry rate	BDS	11.6%
Startup's exit rate	BDS	26.9%
Persistence employment	Sterk et al. (2021)	0.97
Per capita consumption growth	NIPA	2.6%
Persistence growth rates	Dosi et al. (2019)	0.22
St. dev. employment 2-years to 1-year firms	Sterk et al. (2021)	1
Δ growth entrant to incumbent	Huergo and Jaumandreu (2004)	2 pp

entry and exit as in the previous section (to generate a model defined "churning" in the table, setting $f_E = \infty$ and keeping $\lambda = 0$. Finally, I shut down churning, deriving a model defined as "symmetric model" in the table. While holding the number of firms fixed at the level obtained with the no-entry/exit specification, I shut down the R&D productivity shock by setting $\sigma_\xi = 0$ and adjusting $\bar{\alpha}$ to preserve the same mean. In this way, I can compute the aggregate growth rate of the economy under a symmetric equilibrium and compare it to one where the only difference is the presence of churning driven by idiosyncratic shocks to the ability to innovate.

Table 4 summarizes the results. I use the full model as the baseline model. In all other models, the values in red are the arbitrary targets set exogenously.

Table 4: The quantitative importance of turbulence

	Symmetric model	Model with churning	Churning and entry=exit	Full model
Net entry rate	0%	0%	0%	1%
Entry rate	0%	0%	10.5%	11.6%
Firm size	45	45	17	17
Dispersion (productivity 90th to 10th pctl.)	1	1.72	1.34	1.33
Growth	2.9%	1.8%	1.8%	2.6%
Change in R&D to GDP	2700% (+1.82pp)	+180% (+0.12pp)	+1% (+0pp)	0

Removing net entry

Removing net entry by setting population growth to 0 has a noticeable effect on growth and little else. The expanding variety channel makes up approximately 30% of the growth rate, given this value for the elasticity of substitution. Population growth also affects the interest rate, but the results show that this change does not affect the results noticeably.

Removing entry and exit

By comparing the second column with the fourth column, it is clear that entry and exit's contribution to the growth process is not quantitatively important, as only the second decimal of the growth rate is affected. This result is surprising, given how much the literature has emphasized this channel. After all, entry and exit entail the replacement of unproductive and non-innovative firms with others that are, on average, more productive and innovative. Therefore, this selection effect would change the population of firms in the market in a way that promotes faster aggregate productivity growth.

However, this model includes countering endogenous forces. When entry is shut down, the steady state number of firms in the market drops, increasing the average firm size. Consequently, firms can spread the cost of R&D onto more units of product sold, promoting more R&D spending (cost-spreading effect).

Finally, another effect contributes to increasing R&D, although it may not affect growth. When entry does not occur, the market is populated by more low-productivity firms that would otherwise exit. As these firms are small market share, the marginal value of R&D is relatively high. However, as these firms are not good innovators, the impact of their R&D on aggregate growth is minimal. Indeed, it is interesting that productivity growth is unaffected despite higher R&D spending.

The effect on productivity dispersion is instead substantial. In particular, exit plays a vital role in replacing unproductive and non-innovative firms with others that are better able to compete. If exit suddenly stops, these firms stay in the market but start losing productivity relative to their competitors. Ultimately, they reach a size that is low enough to ensure that the forces of attraction prevent them from shrinking further. Meanwhile, the distance between the 90th and 10th percentile of the productivity distribution rises.

All these results suggest that if something were to change, reducing entry and, consequently, exit rates, we should not expect a much different steady-state growth rate through that channel. However, there would be two other relevant effects. First, aggregate R&D would increase without affecting growth. Firms in the market would be less innovative but would invest more. Second, the market would be more concentrated. There would be fewer firms, implying a larger average size, and the size dispersion would increase.

Removing churning

When comparing the models with churning and without churning, the latter produces a higher growth rate by 1.1 percentage points. This outcome is driven by higher R&D investment. With dispersed and skewed R&D productivities, most firms are less innovative than average. Therefore, most firms contribute little to both R&D and productivity growth. Productivity growth is instead disproportionately affected by highly innovative firms. Can these more innovative firms make up for the low contribution of most firms? Only to some extent due to the nonlinearities of the model. First, diminishing returns to R&D imply that as firms increase their investment, they reduce its efficacy. Second, as the most innovative firms grow, they tend to reduce their investment as the market share they can gain with further investment declines. Instead, under symmetry, most firms have a higher R&D productivity than in the model with heterogeneity, leading them to invest more and to create more knowledge per unit of R&D.

The takeaway is that tougher dynamic competition (which is endogenous) is an effect that can spur productivity growth by providing the incentive to conduct more R&D while reducing the dispersion of firm sizes.

6 Summary and Conclusions

This paper presents a unified framework for studying aggregate productivity growth and business dynamism in all dimensions. The model features a continuum of monopolistically competitive firms subject to idiosyncratic shocks to their R&D productivity. Due to in-house R&D and 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 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 and turbulence to the former.

Finally, the quantitative results point out how and why disregarding the presence of entry, exit, and churning can affect the resulting growth rate. In particular, while net entry is relevant for growth, a reduction in entry that is compensated by the same reduction in exit has a negligible effect on growth while increasing considerably R&D spending and firm size dispersion. Instead, a market characterized by firms that are similarly able to innovate displays faster growth and less churning than a market where few innovators drive growth.

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.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018). Innovation, Reallocation, and Growth. *American Economic Review*, 108(11):3450–3491.

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

Acemoglu, D. and Ventura, J. (2002). The World Income Distribution. *The Quarterly Journal of Economics*, 117(2):659–694.

Aghion, P. and Howitt, P. (1990). A Model of Growth Through Creative Destruction. Technical Report w3223, National Bureau of Economic Research, Cambridge, MA.

Akcigit, U. and Ates, S. T. (2019). What Happened to U.S. Business Dynamism? Working Paper 25756, National Bureau of Economic Research.

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

Audretsch, D., Klomp, L., Santarelli, E., and Thurik, A. (2004). Gibrat's Law: Are the Services Different? *Review of Industrial Organization*, 24(3):301–324.

Bilbiie, F. O., Ghironi, F., and Melitz, M. J. (2012). Endogenous Entry, Product Variety, and Business Cycles. *Journal of Political Economy*, 120(2):304–345.

Bloom, N., Schankerman, M., and Reenen, J. V. (2013). Identifying Technology Spillovers and Product Market Rivalry. *Econometrica*, 81(4):1347–1393.

Bottazzi, G., Cefis, E., Dosi, G., and Secchi, A. (2007). Invariances and Diversities in the Patterns of Industrial Evolution: Some Evidence from Italian Manufacturing Industries. *Small Business Economics*, 29(1-2):137–159.

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

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

Caves, R. E. (1998). Industrial Organization and New Findings on the Turnover and Mobility of Firms. *Journal of Economic Literature*, 36(4):1947–1982.

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

Decker, R. A., Haltiwanger, J., Jarmin, R. S., and Miranda, J. (2016). Declining Business Dynamism: What We Know and the Way Forward. *American Economic Review*, 106(5):203–207.

Dinopoulos, E. and Thompson, P. (1998). Schumpeterian Growth Without Scale Effects. *Journal of Economic Growth*, 3(4):313–335.

Dosi, G. (1988). Sources, Procedures, and Microeconomic Effects of Innovation. *Journal of Economic Literature*, 26(3):1120–1171.

Dosi, G., Grazzi, M., Moschella, D., Pisano, G., and Tamagni, F. (2019). Long-term firm growth: An empirical analysis of US manufacturers 1959–2015. *Industrial and Corporate Change*, page dtz044.

Ericson, R. and Pakes, A. (1995). Markov-Perfect Industry Dynamics: A Framework for Empirical Work. *The Review of Economic Studies*, 62(1):53–82.

Foster, L., Haltiwanger, J., and Krizan, C. (2001). Aggregate Productivity Growth: Lessons from Microeconomic Evidence. In *New Developments in Productivity Analysis*, Edited by C.R. Hulten, E.R. Dean, and M.J. Harper, pages 303–363. University of Chicago Press, Chicago.

Garcia-Macia, D., Hsieh, C.-T., and Klenow, P. J. (2019). How Destructive Is Innovation? *Econometrica*, 87(5):1507–1541.

Gomme, P., Ravikumar, B., and Rupert, P. (2011). The return to capital and the business cycle. *Review of Economic Dynamics*, 14(2):262–278.

Ha, J. and 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(4):733–774.

Hall, B. H. and Lerner, J. (2010). Chapter 14 - The Financing of R&D and Innovation. In Hall, B. H. and Rosenberg, N., editors, *Handbook of the Economics of Innovation*, volume 1 of *Handbook of The Economics of Innovation, Vol. 1*, pages 609–639. North-Holland.

Hopenhayn, H. A. (1992). Entry, Exit, and firm Dynamics in Long Run Equilibrium. *Econometrica*, 60(5):1127–1150.

Hopenhayn, H. A. (2014). Firms, Misallocation, and Aggregate Productivity: A Review. *Annual Review of Economics*, 6(1):735–770.

Huergo, E. and Jaumandreu, J. (2004). Firms' age, process innovation and productivity growth. *International Journal of Industrial Organization*, 22(4):541–559.

Klette, T. J. and Kortum, S. (2004). Innovating Firms and Aggregate Innovation. *Journal of Political Economy*, 112(5):986–1018.

Laincz, C. A. (2009). R&D subsidies in a model of growth with dynamic market structure. *Journal of Evolutionary Economics*, 19(5):643–673.

Laincz, C. A. and Peretto, P. F. (2006). Scale effects in endogenous growth theory: An error of aggregation not specification. *Journal of Economic Growth*, 11(3):263–288.

Lee, Y. and Mukoyama, T. (2015). Entry and exit of manufacturing plants over the business cycle. *European Economic Review*, 77:20–27.

Madsen, J. B. (2008). Semi-endogenous versus Schumpeterian growth models: Testing the knowledge production function using international data. *Journal of Economic Growth*, 13(1):1–26.

Mand, M. (2019). On the cyclicality of R&D activities. *Journal of Macroeconomics*, 59:38–58.

Melitz, M. J. and Polanec, S. (2015). Dynamic Olley-Pakes productivity decomposition with entry and exit. *The RAND Journal of Economics*, 46(2):362–375.

Mukoyama, T. and Osotimehin, S. (2019). Barriers to Reallocation and Economic Growth: The Effects of Firing Costs.

Peretto, P. F. (1998). Technological Change, Market Rivalry, and the Evolution of the Capitalist Engine of Growth. *Journal of Economic Growth*, 3(1):53–80.

Peretto, P. F. (2018). Robust endogenous growth. *European Economic Review*, 108:49–77.

Peretto, P. F. and Connolly, M. (2007). The Manhattan Metaphor. *Journal of Economic Growth*, 12(4):329–350.

Peretto, P. F. and Smulders, S. (2002). Technological distance, growth and scale effects. *The Economic Journal*, 112(481):603–624.

Restuccia, D. and Rogerson, R. (2017). The Causes and Costs of Misallocation. *Journal of Economic Perspectives*, 31(3):151–174.

Rossi-Hansberg, E. and Wright, M. L. J. (2007). Establishment Size Dynamics in the Aggregate Economy. *American Economic Review*, 97(5):1639–1666.

Sterk, V., Sedláček, P., and Pugsley, B. (2021). The Nature of Firm Growth. *American Economic Review*, 111(2):547–579.

Sutton, J. (1997). Gibrat's Legacy. *Journal of Economic Literature*, 35(1):40–59.

Thompson, P. (2001). The Microeconomics of an R&D-Based Model of Endogenous Growth. *Journal of Economic Growth*, 6(4):263–283.

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

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:

$$\widetilde{x}_{i,t} = z_{i,t}^\theta l_{x_{i,t}}, \tag{36}$$

where the first order condition for production labor is

$$l_{x_{i,t}} = \left[\frac{\epsilon-1}{\epsilon\widehat{w}_t}\left(\frac{\widehat{y}_t}{n_t}\right)^{\frac{1}{\epsilon}} z_{i,t}^{\theta\frac{\epsilon-1}{\epsilon}}\right]^\epsilon, \tag{37}$$

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

$$\check{p}_{i,t} = \frac{\epsilon}{\epsilon-1}\frac{\widehat{w}_t}{\nu z_{i,t}^\theta}. \tag{38}$$

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

$$\widehat{\pi}_{i,t}(z_{i,t}, l_{Z_{i,t}}) = \widehat{w}_t\left[\frac{\epsilon}{(\epsilon-1)} - 1\right]\left[\frac{\epsilon-1}{\epsilon\widehat{w}_t}\left(\frac{\widehat{y}_t}{n_t}\right)^{\frac{1}{\epsilon}} z_{i,t}^{\theta\frac{\epsilon-1}{\epsilon}}\right]^{\frac{\epsilon}{\epsilon-(\epsilon-1)}} - \widehat{w}_t(l_{Z_{i,t}} + \Phi) \tag{39}$$

The stationary Bellman equation is:

$$\widehat{V}(z_{i,t}) = \max_{\{l_{Z_{i,t}}\}}\left\{\widehat{\pi}_{i,t}(z_{i,t}, l_{Z_{i,t}}) + (1+g_{t+1}^Z)^\theta(1+\lambda)^{\frac{1}{\epsilon-1}}\left(\frac{n_{t+1}}{n_t}\right)^{\frac{1}{\epsilon-1}}\frac{\widehat{c}_{t+1}}{\widehat{c}_t}\times \right.$$
$$\left.\frac{1}{1+r_{t+1}}\max\{\mathbb{E}_t\widehat{V}(z_{i,t+1}), 0\}\right\} \tag{40}$$

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

$$z_{i,t} = \frac{z_{i,t-1} + \alpha_{i,t} z_{i,t-1}^\mu l_{Z_{i,t-1}}^\zeta}{(1+g_t^Z)}, \tag{41}$$

The entry condition (19) is:

$$\mathbb{E}_t\widehat{v^E}_{i,t}(\alpha_{i,t+1}, z_{i,t+1}) \geq f_E, \tag{42}$$

The equilibrium conditions are modified as follows. The labor market clearing (23) becomes:

$$\frac{1}{n_t} = \frac{\int_0^{N_t}(l_{x_{i,t}} + l_{Z_{i,t}})di}{N_t} + \Phi; \tag{43}$$

the law of motion of the number of establishment (20) is now:

$$\frac{n_{t+1}}{n_t} = \frac{1 + \frac{N_{E_t} - N_{X_t}}{N_t}}{1+\lambda}; \tag{44}$$

output (22) is:

$$y_t = c_t + f_E \frac{N_{E_t}}{N_t} n_t; \tag{45}$$

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 (40) 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 (41). 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 (17) and $\alpha_{i,t}$ according to equation (18).

- 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 (44).

- 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 (45) and (43) respectively. Increase the wage if the left side of equation (42) 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.