



Turbulent growth: Business dynamism and aggregate productivity[☆]

Filippo Massari^{ID}

Dolan School of Business, Fairfield University, 1073 North Benson Road, Fairfield, 06824, CT, USA

ARTICLE INFO

Keywords:

Aggregate productivity
Firm dynamics
Turbulence
Endogenous growth
Firm-size distribution

ABSTRACT

Turbulence refers to the endogenous reallocation of resources (such as jobs) across firms due to entry, exit, and churning (movements within the firm-size distribution). The paper develops a model of turbulent endogenous growth in which firms invest in in-house innovation to cut costs and gain market share. As firms grow, the marginal return to market share declines due to downward-sloping demand, weakening the incentive to innovate. This mechanism, combined with idiosyncratic shocks, generates endogenous churning while preserving a stationary firm-size distribution. The results are robust to introducing entry and exit, which amplify churning and affect growth through selection. In a counterfactual exercise, I model the observed decline in high-growth startups as a thinning of the right tail of the R&D productivity distribution. While eliminating skewness can generate large reductions in aggregate outcomes, matching its decline explains only 15% of the post-2000 slowdown, suggesting a limited aggregate role for fast-growing startups.

1. Introduction

Aggregate productivity growth is central to improving living standards and a core focus of modern growth theory. A large theoretical and empirical literature emphasizes firm-level R&D as the primary driver of productivity growth, linking innovation to firm growth and entry, and its absence to firm exit. As a result, aggregate productivity growth is inherently a *turbulent process*, intrinsically connected to business dynamism.¹

In this context, the observed decline in U.S. rates of entry, exit, and churning – mirrored in other high-income economies – has raised concerns that the same causes behind this decline could explain the productivity growth slowdown, which occurred around the same time.² However, studying aggregate productivity growth and turbulence jointly poses a theoretical challenge: while constant returns to knowledge are needed for growth, they can imply a counterfactual tendency to monopoly. A proper investigation of the phenomenon would have to carefully characterize the process that delivers a stationary firm-size distribution.

[☆] I am grateful to my advisors Pietro Peretto, Giuseppe Fiori, and James Nason for their guidance and encouragement. I thank Domenico Ferraro, Daisoon Kim, Chris Laincz, Soroush Ghazi, Ayse Dur, Robert Kane, the editor Evi Pappa, the associate editor, referees, and participants to seminars at North Carolina State University, Fairfield University, the University of Milan, and Drexel University. Additionally, part of the work was supported by the Jenkins Fellowship at NC State University. Lastly, I acknowledge the support of the High-Performance Computing Services at NC State University. All errors are my own.

E-mail address: fmassari@fairfield.edu.

URL: <https://www.filippomassari.com/>.

¹ Turbulence refers to changes in firm demographics through entry, exit, and churning — the reshuffling of firms' market shares as they grow or shrink. Brown et al. (2008, p. 3) define it as “the entire process of economic change: worker reallocation as workers change jobs and job reallocation from firms contracting and shutting down, to firms expanding and starting up”.

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

This paper develops a rich yet parsimonious framework to study the relationship between turbulence and growth when firms differ in their ability to innovate. The model features a stationary firm-size distribution with endogenous churning. Stationarity arises because more innovative firms grow faster when small, but as they expand their market share, the return to further innovation declines. This mechanism defines a *no-churning locus*: combinations of firm size and innovation ability at which firms grow at the aggregate rate and maintain their relative position. Firms away from the locus adjust their R&D and market share as they converge to it, generating churning in the process. Because R&D also drives long-run growth, the framework links turbulence directly to aggregate productivity growth. This structure highlights the central role of firms that operate far from the no-churning locus in driving both turbulence and growth.

Which firms are most responsible for turbulence? Those hit by large shocks that displace them from the no-churning combination of size and innovation ability, and entrants that begin far from it. The latter includes highly innovative small startups, whose disappearance in the U.S. has been well documented (Decker et al., 2016b). I embed the disappearance of high-growth startups into this framework to assess quantitatively the connection between turbulence and growth, modeling it as a thinning of the right tail of entrants' innovation ability. Because these firms start small but grow quickly, they operate far from the no-churning locus and disproportionately contribute to both turbulence and aggregate growth as they scale.

The theoretical mechanism rests on three standard assumptions: (i) in-house R&D; (ii) firm-specific idiosyncratic shocks – in this case to the R&D productivity – drawn from a common distribution, which drive heterogeneity in knowledge stocks and market shares; and (iii) imperfect substitutability between goods, which leads to diminishing returns to *relative* knowledge. The incentive to innovate arises from the potential to gain market share. Yet this same force gives rise to firm-level diminishing returns: as firms expand output, they must reduce prices, lowering the return to further market share gains. As a result, all else equal, the incentive to innovate declines with product market share, generating a negative relationship between growth and firm size.³

Do firm-level economic diminishing returns imply that aggregate productivity growth will eventually decline to zero? No. Because these diminishing returns apply in relative terms, it is still possible to sustain increasing returns in absolute terms. The key condition for endogenous growth is that returns to innovation are constant on average — that is, the distribution of returns is stationary. Under this condition, R&D investment and firm growth rates remain constant on average, yielding a constant aggregate growth rate. Diminishing returns at the firm level imply a stationary distribution of market shares, which supports a stationary distribution of innovation returns and, in turn, a stationary and ergodic distribution of firm growth rates. As in other fully endogenous growth models, aggregate growth depends on firms' optimal R&D decisions.

Allowing for endogenous, simultaneous entry and exit does not disturb this process but adds several appealing features. First, it improves realism, reflecting the high turnover observed in most industries. Second, it introduces selection effects that shape firms' life cycles. Third, it affects the aggregate growth rate through its impact on the average R&D productivity in the economy. Finally, it influences churning, as entrants typically begin far from the no-churning locus.

In a quantitative exercise, I model the post-2000 disappearance of high-growth startups as a thinning of the right tail of entrants' innovation ability. At the magnitude observed in the data, and despite the prominent role of these firms in the model, this mechanism accounts for roughly 15% of the decline in productivity growth and for a comparable share of the decline in entry and job reallocation among incumbents. Motivated by this finding, I then consider, as a complementary exercise, a broader decline in innovative abilities affecting both incumbents and entrants symmetrically, calibrated to match the observed slowdown in productivity growth. I then ask whether such a mechanism can also account for a much larger fraction of the observed decline in turbulence.

Literature Review. This paper builds on the firm dynamics and endogenous growth literatures — particularly (Hopenhayn, 1992; Peretto and Connolly, 2007). Central to this framework is the concept of a no-churning locus: an endogenous combination of firm size and ability to innovate at which firms' productivity grows at the aggregate rate. Firms dynamically converge towards it, and their deviation from it drives churning and growth heterogeneity.

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

The model is also closely related to the industrial-organization-based growth framework developed by Peretto and Connolly (2007), which builds on Peretto (1999) and whose theoretical variants and applications are presented in Peretto (2026). As in that tradition, it features both vertical innovation — cost reduction in the production of existing goods — and horizontal innovation — the introduction of new products — where vertical innovation drives long-run growth and horizontal innovation determines the equilibrium number of varieties and the intensity of product-level competition. This paper extends the Peretto framework

³ Although the paper focuses on the product line – as the relevant unit for aggregate productivity growth – I use the term firm synonymously throughout, since each firm produces a single product. Similarly, entry refers to horizontal innovation, i.e., the introduction of a new product.

by introducing firm-specific R&D productivity shocks, which generate endogenous churning, simultaneous entry and exit, and a non-degenerate firm size distribution. Embedding these features in a Peretto-style environment is particularly valuable given the framework's established usefulness for addressing industrial organization questions within a growth context, and allows the model to speak directly to turbulence alongside growth.

Endogenous growth models with heterogeneous producers are not new; their intellectual foundations lie in industrial organization, notably (Ericson and Pakes, 1995). Closely related contributions include (Thompson, 2001; Laincz, 2009), which study growth under heterogeneous firms and dynamic product-level competition. Relative to this paper, Thompson (2001) abstracts from economic diminishing returns to endogenous productivity by assuming that R&D incentives are independent of firm size. In contrast, in my model the dependence of R&D incentives on firm size is a central driver of turbulence, generating a feedback from the firm-size distribution to innovation decisions and aggregate growth. Laincz (2009) instead features a tendency towards monopoly counteracted by technological diffusion from the leader to entrants, whereas my model delivers a non-degenerate firm-size distribution through product differentiation. The two frameworks are therefore complementary and describe distinct market environments.

Much of the recent literature on firm dynamics and growth builds on the Klette and Kortum (2004) framework but typically relies on specific assumptions about entry and exit to ensure a stationary firm size distribution.⁴ In contrast, this paper derives stationarity from firms' endogenous R&D decisions, through a no-churning locus that acts as an attractor in the state space defined by firm size and innovation ability. This mechanism – grounded in standard industrial organization principles – links turbulence and growth directly through the firm's optimization problem.

The quantitative exercise of this paper is related to Olmstead-Rumsey (2020), which studies the link between declining growth and rising concentration by reducing the innovation advantage of laggard firms. The two papers differ in both structure and focus. Olmstead-Rumsey (2020) builds on the quality-ladder tradition in a duopoly with a competitive fringe environment and studies concentration dynamics measured by the market share of the largest firm. By contrast, this paper adopts a Hopenhayn-style framework with monopolistic competition, heterogeneous entrants, and endogenous exit, which is better suited to studying turbulence arising from entry, exit, and firm-level churning. In this setting, innovation ability is distributed across firms, and a firm's growth depends jointly on size and innovative capacity, generating a two-dimensional no-churning locus. This structure allows the analysis to discipline higher-order moments of firm dynamics, such as skewness in the growth rate of entrants. While the papers share an interest in how innovation incentives shape aggregate outcomes, their mechanisms are complementary: reduced innovation by laggards primarily affects concentration, whereas diminished innovative capacity among entrants is particularly relevant for turbulence.

2. A model of turbulent growth

The model features discrete time and a monopolistically competitive intermediate sector with a mass of firms, each producing a unique good sold to a perfectly competitive final sector. The final sector aggregates these goods into a single output, used for household consumption and firm creation. Innovation occurs along two dimensions: *technological depth*, via process innovation for existing goods, and *technological breadth*, via the introduction of new goods. Firms face idiosyncratic shocks to R&D productivity and decide whether to exit after observing their shock. Those that stay hire labor for production and R&D to lower future costs. New firms pay a sunk entry cost in output units to introduce a new good and decide whether to exit after their first draw. Aggregate variables evolve deterministically. Households supply labor inelastically and choose consumption and saving.

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,

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

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:

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

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

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

⁴ Relevant works include (Luttmer, 2007; Lentz and Mortensen, 2008, 2016; Acemoglu and Cao, 2015; Acemoglu et al., 2018; Akcigit and Kerr, 2018; Peters, 2020; De Ridder, 2024).

2.2. Final sector

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 its 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

The intermediate sector, populated by N_t firms producing a unique good, consists of incumbents, entrants, and exiting firms. The demand schedule derived above implies a revenue per good of:

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

which can be decomposed into market size and market share.⁵ 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 Eq. (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.

The firm's problem has a static and a dynamic component. I separate them 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.

Static problem: Dividends

In each period, dividends are given by

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

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

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

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.

⁵ By rearranging Eq. (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.

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

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

A labor demand schedule as in (10) implies pricing at a constant markup:

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

Importantly, firms can reduce their relative price by improving their technological knowledge.

As anticipated earlier, from Eq. (7), reducing the relative price, thus gaining market share, is the only way for firms to increase their revenue. Therefore, Eq. (11) illustrates the fundamental way in which dynamic 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 Eq. (8) and using Eq. (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 takes an investment decision to increase its future knowledge, thus reducing its production cost. The R&D productivity is subject to an idiosyncratic shock, driving heterogeneity in firms' productivity level.

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

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

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

An R&D technology of this kind captures four important elements. First, new knowledge is a function of the existing stock of knowledge due to its cumulative nature, i.e. new knowledge builds on existing knowledge. The linearity is the simplest and most tractable specification in which knowledge is the factor that drives long-run exponential growth at a constant rate.⁶

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:

$$\log \alpha_{i,t} = (1 - \rho) \log \bar{\alpha} + \rho \log \alpha_{i,t-1} + \xi_{i,t}, \quad \xi_{i,t} \sim N(0, \sigma_{\alpha}) \tag{13}$$

$0 \leq \rho < 1$

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

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

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

Future profits are discounted using the risk-free interest rate r , which is determined by the representative household's time preferences. The dynamic optimization can be re-expressed as a Bellman equation:

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

constrained by the knowledge accumulation Eq. (12).

⁶ Peretto (2018) and Massari and Peretto (2026) provide a generalization that allows for new knowledge to exhibit increasing or decreasing returns in the existing stock of knowledge.

Lastly, as captured by the max operator, firms face an exit decision at the beginning 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. Firms set up at time t face the same problem as incumbents in period $t + 1$.

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

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

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

$$\alpha_{t+1}^d \sim \text{Lognormal}(\chi_\alpha, \sigma_\alpha^E). \quad (17)$$

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

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:

$$v_t^E = \mathbb{E}_t V(\alpha_{t+1}^d, Z_{t+1}^d) \geq Z_t^\theta N_t^{\frac{1}{c-1}} f_E, \quad (18)$$

where the right side of the inequality is the entry cost made up of a fixed component f_E and of the technological depth, Z_t^θ , and breadth, $N_t^{\frac{1}{c-1}}$, 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 is 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.

2.4. Equilibrium

The equilibrium of the model is defined by:

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

$$N_{t+1} = N_t + N_{E_t} - N_{X_t}, \quad (19)$$

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

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

such that the following conditions hold.

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

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

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

$$Y_t = c_t L_t + Z_t^\theta N_t^{\frac{1}{c-1}} f_E N_{E_t}. \quad (21)$$

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

$$L_t = \int_0^{N_t} (l_{x_{i,t}} + l_{z_{i,t}}) di + \phi N_t. \quad (22)$$

2.4.1. Steady-state

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

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

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

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

and the number of firms per capita

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

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

Furthermore, to simplify the notation, define productivity as:

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

Aggregate productivity level

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

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

Aggregate productivity

where $S_{i,t} = l_{x_{i,t}}/L_{x_t}$ is the production labor share for firm i .⁷ 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:

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

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

Finally, the last term shows that aggregate productivity depends on the distribution of weighted firm relative productivities, as $a_{i,t} = A_{i,t}/A_t$. This element describes the allocative efficiency of the economy. The distribution of individual productivities matters for aggregate productivity because the CES aggregator is a power mean, which is altered by the firms’ relative productivity distribution. To understand why this term is tied to the distribution of productivities and labor share, it is useful to notice that Eq. (27) can be re-expressed as:

$$\frac{1}{N_t} \int_0^{N_t} S_{i,t} N_t a_{i,t} di = 1. \tag{28}$$

⁷ I obtain Eq. (26) by plugging the production function (9) into the CES aggregator (4), then multiplying and dividing by $L_{x_t} A_t$. A_t at the denominator moves into the parenthesis to divide $A_{i,t}$. L_{x_t} moves into the parenthesis to divide $l_{x_{i,t}}$. Next, to isolate the relevance of N_t , I multiply and divide by N_t the factors in the parenthesis. I bring the denominator out, and break it into two parts: one that remains within the square bracket, and the other that comes out denoting the technological breadth (love of variety effect).

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

Aggregate productivity growth rate

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

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

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

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

$$1 + g^{productivity} = \underbrace{\left[\frac{n_t(z_t, \alpha_t)}{n_{t-1}(z_{t-1}, \alpha_{t-1})} (1 + \lambda) \right]^{\frac{1}{\epsilon-1}}}_{\text{semi-endogenous}} \underbrace{(1 + g_t^A(z_{t-1}, \alpha_{t-1}))}_{\text{average productivity}} \underbrace{\left\{ \frac{\frac{1}{N_t} \int_0^{N_t} (S_{i,t} N_t a_{i,t})^{\frac{\epsilon-1}{\epsilon}} di}{\frac{1}{N_{t-1}} \int_0^{N_{t-1}} (S_{i,t-1} N_{t-1} a_{i,t-1})^{\frac{\epsilon-1}{\epsilon}} di} \right\}^{\frac{\epsilon}{\epsilon-1}}}_{\text{change in the distribution}}, \tag{29}$$

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

3. Sources of firm growth, churning, and stationarity

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

This interdependence raises a key question: how can there be (i) growth rate differentials, (ii) constant returns to the growth driving factor (and increasing returns to the private factors overall), and (iii) a stationary firm size distribution? The elements (i) and (ii) may suggest a higher growth rate for relatively larger firms, thus promoting a tendency towards monopoly and violating element (iii).

To better illustrate the mechanism, I temporarily shut down entry and exit. I will reintroduce them in the next section where I discuss their role in detail. I do that by setting the parameters $\lambda = 0$, $f_E = \infty$, $\phi = 0$. The next two subsections discuss the mechanism that preserves the firm size distribution stationarity and its implications.

3.1. Growth rate differentials without entry and exit

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

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

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

$$g_{t+1}^{A_i} \approx \alpha_{i,t} z_{i,t}^{\mu-1} l_{Z_{i,t}}(z_{i,t}, \mathbb{E}_t \alpha_{i,t+1})^\zeta. \tag{30}$$

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

Second, R&D is a function of relative knowledge. The value maximization yields the following policy rule:

$$l_{Z_{i,t}}^{1-\zeta} = \frac{1}{1+r_{t+1}} \frac{\alpha_{i,t} z_{i,t}^\mu}{w_t (1+g_{t+1}^Z)} \mathbb{E}_t \left[\frac{w_{t+1} l_{Z_{i,t+1}}^{1-\zeta}}{\alpha_{i,t+1} z_{i,t+1}^\mu} + \zeta \frac{\partial \pi_{i,t+1}}{\partial z_{i,t+1}} + \frac{\mu w_{t+1} l_{Z_{i,t+1}}}{z_{i,t+1}} \right], \tag{31}$$

with

$$\frac{\partial \pi_{i,t}}{\partial z_{i,t}} = \theta(\epsilon - 1) w_t \left[\frac{\epsilon}{(\epsilon - 1)} - 1 \right] \left[Z^{\theta \frac{\epsilon-1}{\epsilon}} \frac{\epsilon - 1}{\epsilon w_t} \left(\frac{y_t}{n_t} \right)^{\frac{1}{\epsilon}} z_{i,t}^{\theta \frac{\epsilon-1}{\epsilon} - 1} \right]^\epsilon. \tag{32}$$

Eq. (31) shows that firms will choose R&D investment by balancing the present value of relative knowledge’s marginal benefit and marginal cost. The term outside the bracket is the inverse of the marginal cost of new relative knowledge — with a slight modification as I have kept the diminishing returns to R&D on the left side. It increases with the price of R&D and with average knowledge growth, as faster average knowledge growth requires more investment for firms to keep up with the others. It instead decreases with R&D productivity and relative knowledge to the extent that firms internalize it, as these two elements determine the efficacy of R&D.

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

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

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

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

In particular, μ 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, θ , compounds this effect by mapping differences in knowledge levels into differences in firm size, production and, ultimately, profits. Therefore, the parameters θ and μ are essential in determining the shape of the firm size distribution and raise concerns about its non-degeneracy.

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 paper emphasizes this aspect because it implies that standard modeling assumptions deliver stationarity and endogenous churning. They do so by creating an endogenous combination of ability to innovate and size that is the attractor of the endogenous state variable. The next subsection illustrates how this result arises. I focus on the case in which the forces of convergence prevail. The existing literature has addressed all other cases, as I illustrate later.

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 firms’ 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.




Fig. 1. Turbulence, firm growth, and stability.

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

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

The convergence process over relative productivity can be understood by analyzing the R&D technology given in Eq. (12), combined with the policy function (31), which can be rearranged to yield:

$$\frac{a_{i,t+1}(z_{i,t+1})}{a_{i,t}(z_{i,t})} = \frac{(1 + \alpha_{i,t} z_{i,t}^{\mu-1} l_{Z_{i,t}}(\alpha_{i,t}, z_{i,t})^\zeta)^\theta}{1 + g_{t+1}^A} \tag{33}$$

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:

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

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

The conceptualization of a no-churning locus and its implications are the key contributions of this paper. Importantly, the mechanism that delivers endogenous churning and a stationary firm size distribution is responsible for deviating from Gibrat Law, i.e. the hypothesized absence of any correlation of firm growth and relative size. 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.

Because of its tractability, much of the literature relies on assumptions that deliver Gibrat Law (Klette and Kortum, 2004; Acemoglu and Cao, 2015; Acemoglu et al., 2018). Some models depart from Gibrat Law, but through mechanisms different from those emphasized here.⁸

The framework can also generate deviations from Gibrat Law in the opposite direction, with larger firms growing faster and a tendency towards monopoly. In the limit, such parameterizations resemble models of creative destruction with a monopolist and an innovating entrant, such as Aghion and Howitt (1992), or highly concentrated industries as in Laincz (2009). In contrast to those settings, however, the present framework nests both convergent and divergent firm dynamics within a unified structure, depending on the relative strength of innovation incentives and competitive forces.

Estimates of Gibrat's law coefficients – such as those reported in Bottazzi et al. (2007) – vary across industries, suggesting that strict proportional growth is not a general empirical regularity. While some sectors may exhibit tendencies towards monopoly, this is not the norm across most industries. The relevance of productivity churning for aggregate growth can be assessed using standard decompositions, such as Foster et al. (2001) or Melitz and Polanec (2015), incorporating the modification proposed in Massari (2025). These decompositions quantify the direct contribution of churning to aggregate productivity growth. However, deviations from Gibrat Law also affect firms' investment incentives, influencing growth dynamics through channels that such accounting exercises do not capture.

3.3. Stationarity in a simplified environment

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

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

- $\alpha_{i,t}$ is fixed, strictly positive, and finite;
- Agents are sufficiently impatient to ensure that discounting by two periods yields approximately 0, while discounting by one period yields positive values.

the firm size distribution is stationary under two requirements: $\zeta\theta(\epsilon - 1) < 1 - \mu$, and $[\theta(\epsilon - 1) - 1]G < 1$, with G being a positive combination of variables and parameters.

The first assumption helps by removing the expectation operator from the policy function. Proving stationarity under fixed differences is particularly noteworthy as it highlights one of the strengths of this model relative to the existing literature: stationarity of the firm size distribution does not arise out of assuming that differences across firms fade away with time. Instead, it arises because of an endogenous market mechanism arising from standard industrial organization assumptions that disciplines firms' optimal investment decisions. The second assumption is motivated by algebraic convenience, as it simplifies the differentiation of the policy function with respect to relative knowledge.

Of the two requirements that ensure stationarity, the first is the more interesting. It highlights the relationship between the forces of convergence and divergence that determine whether the distribution converges, diverges, or is described by Gibrat Law. The second requirement arises from the discrete nature of the problem and vanishes when the problem is specified in continuous time. Moreover, it is always satisfied when profit is concave in relative knowledge.

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

4. Entry and exit

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

What ensures that the presence of entry and exit will preserve the results illustrated above regarding the stationarity and non-degeneracy of the firm-size distribution? The model has one firm-specific state variable, the knowledge stock, and the firm-specific exogenous shock. The exogenous shock is stationary by assumption. Furthermore, under the set of parameter values under consideration, the analysis provided in the previous section suggests – by proving it in a simplified environment – that relative

⁸ For example, in Akgicig and Kerr (2018) deviations arise because incumbents engage in horizontal innovation whose intensity decreases with firm size, while vertical innovation – arguably the main driver of growth (Garcia-Macia et al., 2019) – is independent of size. The two frameworks therefore offer complementary explanations for deviations from Gibrat Law. Similarly, Thompson (2001) generates declining growth rates with size even when R&D is independent of it; my model can reproduce this outcome under specific parameterizations, though doing so removes economically relevant forces.

knowledge evolves according to a stationary process. To the extent that this last claim is true, the conditions used in [Hopenhayn \(1992\)](#) to prove Theorem 3 on 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.

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 per capita remains constant over time.

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

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

4.1. Effects linking entry, exit, and growth

Entry, exit, and growth interact through several effects. In this subsection, I illustrate them.

Effect 1: Expanding variety

Assuming a CES form for the aggregation of different goods implies a love for variety. This effect matters for growth as shown in (26), where the output growth rate depends on the technological breadth of production.

Effect 2: Level replacement

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

Effect 3: Growth replacement

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

Effect 4: Cost spreading

A parameter change that facilitates entry, will increase the steady state number of firms. Conversely, a parameter change that facilitates exit will decrease the steady state number of firms. A change in the number of firms is a change in the average market share, which has consequences for R&D investment. From Eq. (31), R&D investment decreases in the number of firms. The cost spreading effect explains this result: as the cost of innovation is spread on all units sold, a more crowded market where each firm sells fewer units reduces the incentives to innovate. This effect is emphasized in the symmetric version of this model ([Peretto and Connolly, 2007](#)).

Effect 5: R&D-growth decoupling

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

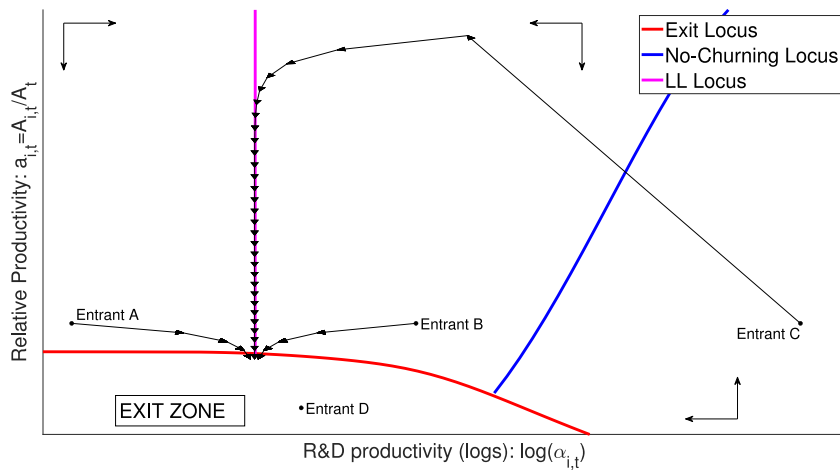


Fig. 2. The firm expected life cycle.

Note: Phase diagram over the state variable and the exogenous shock. The exit locus shows the combination of relative productivity and R&D productivity below which the firm exits the market. The other curves are those introduced in Fig. 1. The figure also shows the expected life path of four startups that differ in their initial draws.

4.2. Prediction 2: Firm life cycle

As in Hopenhayn, the model provides predictions over firms' life cycle. However, as a firm's productivity is not a random draw, the life cycle differs. Fig. 2 re-proposes the phase diagram of Section 3.2 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 once they reach it. 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.

Endogenous entry and exit determines all the effects discussed in the previous subsection. Firm life cycle dynamics are shaped by one of them, the growth replacement effect, which ensures that the R&D productivity cross-sectional average is higher than its unconditional expectation depicted on the LL locus. Consequently, firms expect to converge to the exit locus in finite time, as shown in Fig. 2.

Furthermore, the phase diagram highlights 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 four arbitrary entrants depicted in the diagram illustrate this point.

Some (Entrants A and B) exit early due to poor innovation draws; others (Entrant D) exit immediately due to poor draws overall. 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 innovates it reaches 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. Absent any new 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) drives an imperfect substitute (for example, a VHS player producer) out of the market over time by gradually increasing its relative efficiency.

Since the quantitative exercise proposed in this model concerns the disappearance of high R&D productivity entrants, what does the model say about their effect on churning? Because entrants start small on average, high R&D productivity entrants will cluster in the bottom right corner of Fig. 2. Therefore, they will initially be far from the no-churning locus. As a result, they will grow fast to reach it, thus contributing disproportionately to churning and aggregate growth.

Finally, the figure shows the intersection between the no-churning locus and the LL locus within the exit zone. In this case, firms eventually exit the market with probability 1. This is the more realistic scenario and the one that emerges under the calibration presented below. However, a different set of parameters leads the LL locus to intersect the no-churning locus outside the exit zone. Of particular interest is the case where the exit locus and the no-churning locus never intersect. This implies that there exists a level $z_{i,t} = z^*$ such that $v(z^*, 0) > 0$ and $g_i^A(z^*, \alpha_{i,t}) > g^A$ for any $\alpha_{i,t}$. Under this parametrization, any firm that reaches size z^* survives indefinitely. Thus, the firm life cycle dynamics differ from those shown in the figure, but the result on the stationarity of the firm size distribution is preserved, since the growth-limiting mechanism described in the previous section ensures stationarity without requiring firm exit.

Table 1
Externally calibrated parameters.

Parameter & target	Symbol & value
Labor force growth	$\lambda = 1.9\%$
Returns to R&D labor	$\zeta = 0.55$
Discount rate	$\beta = 0.98$
Elasticity of substitution between goods	$\epsilon = 3.9$
Persistence R&D productivity	$\rho = 0.71$
Standard deviation R&D productivity shock	$\sigma_\alpha = 1.89$
Returns to knowledge in production	$\theta = 0.1$
Private vs. social knowledge	$\mu = 0.33$

5. Quantitative exercise: The disappearance of innovative startups

This section evaluates whether the decline in highly innovative startups can account for the observed reductions in turbulence and productivity growth since the 2000s.

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

Meanwhile, the aggregate productivity growth rate declined significantly from the mid-2000s onward, relative to the previous decade, returning to levels similar to those observed from the mid-1970s to the mid-1990s (Byrne et al., 2016; Syverson, 2017; Fernald, 2018).

To answer the question, I conduct a comparative statics exercise by comparing model steady states under the assumption that the only change is a thinning of the right tail of the distribution of entrants' ability to innovate, in order to match the documented decline of high-growth startups in the U.S. (Decker et al., 2016b).

I choose 2003 as the break year because it corresponds to the largest change in the job reallocation rate among incumbents.⁹ Consistent with the model's timing assumption that changes in growth respond with a one-period lag to parameter changes, this implies 2004 as the break year for productivity growth, in line with the existing literature.

5.1. Calibration

This subsection presents the calibration, which consists in matching selected moments. I pick as many moments as parameters to calibrate, so that I can match them exactly. As data on business dynamism are available from 1978 to 2019, I rely on averages at an annual frequency from 1978 to 2003, unless stated otherwise. I divide my discussion in externally calibrated parameters, namely those that have a one-to-one correspondence with a selected moment, and internally calibrated parameters, those that interact with each other to deliver the targeted moment.

Externally calibrated parameters

The externally calibrated parameters are summarized in Table 1. The employment growth rate, λ , for the U.S. averages 1.9% per year. The parameter ζ is set to 0.55 to match the return to labor in R&D in the U.S. based on NSF data (Mand, 2019). $\beta = 0.98$ is selected to match, anticipating a growth rate of per capita consumption of approximately 2%, a real rate of return of approximately 4% (Gomme et al., 2011). I further set $\epsilon = 3.9$ to match a markup over marginal cost of 35% (De Ridder et al., 2022).

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

Finally, I estimate from Compustat data the parameters ρ and σ_α . I use all firms with positive sales and R&D values who have at least three consecutive observations (which is the minimum number of consecutive periods needed to pursue the estimation) from 1978 to 2019.¹⁰ From the computed values of $\alpha_{i,t}$, I estimate the AR(1) process of Eq. (13), finding $\rho = 0.71$ and $\sigma_\alpha = 1.89$.

⁹ In the data, the job reallocation rate is defined as the sum of job creation and job destruction rates divided by average employment, using creation and destruction at continuing establishments between years $t-1$ and t . In the model, I compute the same statistic using the simulated model with one hundred thousand establishments.

¹⁰ I compute their level of technological knowledge using Eq. (9). In that equation, I divide sales by employment to find productivity which in the model corresponds to $Z_{i,t}^\theta$. Then, I use Eq. (12) while assuming that the relevant parameters take the values presented in this section. Specifically, I raise the productivity computed earlier to the power $1/\theta$ to obtain knowledge. I then solve the equation for the $\alpha_{i,t}$ of each firm at each time period by using the $Z_{i,t}$ just computed, R&D data, and the values of μ and ζ mentioned in the previous paragraphs.

Internally calibrated parameters

The remaining parameters, summarized in [Table 2](#), jointly determine the targeted moments implied by the model. Each targeted moment disciplines a distinct margin of the stationary equilibrium (see [Appendix A](#) for the stationary equilibrium conditions): the entry margin (f_E), the operating-cost/scale margin (ϕ), the level and dispersion of entrants' relative productivity at entry (χ_Z, σ_Z^E), the mean and tail asymmetry of entrants' innovation outcomes ($\chi_\alpha, \sigma_\alpha^E$), and the aggregate growth rate ($\bar{\alpha}$). For what concerns entrants, the model would ideally require product-level data. Due to data availability issues, I rely on establishment-level data.

Entry and scale. The fixed entry cost f_E is chosen to match the entry rate of new establishments which, according to U.S. Census Business Dynamics Statistics (BDS), averages 12.7%.

The fixed operating cost ϕ is chosen to match average establishment size, $1/n = L/N$, which equals 16 workers in the BDS data. In the model, average establishment size is an equilibrium outcome, jointly determined with entry, exit, and labor allocation. This scale margin is central in this class of models, as establishment size determines the extent of cost spreading across firms and thereby affects incentives for R&D.

The logic underlying average establishment size as an identifying restriction is as follows. Targeting average establishment size pins down the average amount of labor available per firm. Since overhead labor demand is fixed per establishment, this pins down the fraction of firm-level labor absorbed by overhead for a given fixed operating cost ϕ .¹¹ As a result, once average establishment size is fixed, any increase in ϕ mechanically crowds out labor available for production and R&D via the labor market clearing condition, thereby affecting profits, entry, innovation, and aggregate growth.

This logic can be made precise by considering the labor market clearing condition expressed at the level of the average firm in Eq. (A.8). Holding the remaining structural parameters fixed pins down firms' production and R&D labor demands through their first order conditions. Targeting average establishment size therefore fixes the labor supply side of this equation, that is, the total amount of labor available per firm. Since overhead labor enters additively and is fixed per establishment, any change in ϕ shifts the labor demand side. There is no direct firm-level margin through which this change can be absorbed while holding labor per firm fixed: adjustments in production or R&D labor alter profits and firm values, thereby affecting entry and exit and changing n . Consequently, holding the other parameters fixed, the labor market clearing condition implies a mapping between ϕ and n .

Among empirically meaningful scale targets, I use average employment per establishment to discipline the fixed operating cost ϕ . The main reason is that this moment is directly connected to the average firm labor market clearing condition, in which average establishment size pins down the amount of labor available per firm and ϕ enters additively as overhead labor. This makes it not only a natural scale-based restriction on ϕ , but also a convenient target from a computational standpoint, since it can be matched directly through that equilibrium condition. A second reason is empirical: average employment per establishment is measured from Census establishment counts and employment totals, making it a relatively transparent and reliable target. Finally, this choice is also supported by the model's untargeted implications. As discussed below, calibrating the model to match average establishment size yields scale-related moments – such as the ratio of production to total employment and the private R&D intensity of the economy – that remain reasonably close to their empirical counterparts.

Entrants' productivity level and dispersion. Following [Lee and Mukoyama \(2015\)](#), the average relative productivity of entrants is 0.96, which disciplines the mean entrant initial knowledge parameter χ_Z . I calibrate the dispersion of entrants' initial productivity, σ_Z^E , to match the exit rate of one-year-old establishments from the BDS, which is 26.9%. This target identifies σ_Z^E because exit in the model is driven by the lower tail of the entrant productivity distribution, holding fixed the mean level of entrant productivity.

Entrants' innovation outcomes. Next, I choose χ_α and σ_α^E , which govern the distribution of entrants' R&D productivity. The mean entrant R&D productivity parameter χ_α is chosen to match evidence on the contribution of young firms to productivity growth. Specifically, I rely on [Foster et al. \(2008\)](#)'s productivity growth accounting decomposition at a five-year horizon and match the share of productivity growth attributed to entry, 24%. I use ([Foster et al., 2008](#)) because it covers more years of observation and uses improved techniques for isolating productivity relative to earlier alternatives ([Foster et al., 2001](#)). To discipline the dispersion and asymmetry of entrants' innovation outcomes, I choose σ_α^E to match the difference between the 90th-50th and 50th-10th percentiles of the startup growth distribution reported by [Decker et al. \(2016b\)](#), which is a measure of skewness.

Aggregate growth. Finally, I choose $\bar{\alpha}$ to match an average labor productivity growth rate of 2.1%. Importantly, following [Bilbiie et al. \(2012\)](#), I distinguish between data-consistent moments and model-consistent moments when it comes to growth. Note that in this model labor productivity grows because of an increase in average productivity and because of a variety expansion effect. Since the data collection process misses the effect on productivity through variety expansion, the target of 2.1% is for average labor productivity growth (see [Table 2](#)).

Untargeted moments

A few untargeted moments deserve attention. First, since the objective of this model is to understand the consequences of a parameter change for turbulence and growth, the only relevant moment in this regard that is not directly targeted in the calibration is the job reallocation rate. The calibrated model accounts for 41% of the observed job reallocation rate among incumbents. This is a non-trivial share of overall reallocation, especially considering that productivity differences across firms are influenced by many

¹¹ Mathematically, this can be seen by dividing aggregate overhead labor ϕn by total labor L , which implies that the overhead labor share equals ϕn , where $n \equiv N/L$.

Table 2
Internally calibrated parameters.

Parameter	Symbol & value	Main target
Entrant's mean R&D productivity	$\log \chi_a = 1.09$	Entrants' growth contribution: 24%
Entrant's mean initial knowledge	$\log \chi_z = -3.88$	Average productivity of entrants: 0.96
Standard deviation entrants' knowledge	$\sigma_z^E = 2.20$	Entrants' exit rate: 26.9%
Standard deviation entrants' R&D productivity	$\sigma_a^E = 1.22$	90 to 50 pctlte - 50 to 10 pctlte: 0.17
Fixed operating cost	$\phi = 2.99$	Average establishment size: 16
Fixed entry cost	$f_E = 5.23$	Establishment entry rate: 12.7%
R&D productivity	$\log \bar{\alpha} = -7.42$	Labor productivity growth: 2.1%




Fig. 3. The firm-size distribution.

Note: Comparison between a counter cumulative distribution function (in logarithmic scale) for the firm sizes generated by the model (10,000 bins from 100,000 observations) and for a fitted Pareto distribution and lognormal distribution, estimated via maximum likelihood from the model-generated data.

other factors – such as the external environment – which are often subject to change (Syverson, 2011). For the same reason, the model account only for a similar non-trivial share of the dispersion in the productivity distribution across firms. Compared to an average across industries and years 1997 to 2016 computed on BLS data, the model accounts for approximately 36% of the dispersion in total factor productivity.

Second, because this is a growth model, closely matching the share of resources devoted to innovation is important. The model generates a private R&D-to-GDP ratio of 2.2%, which is close to the 1.5% observed in the data for the period 1978–2003, and even closer to the more appropriate ratio of 1.9% where government expenditure is excluded from GDP.

Additionally, Fig. 3 plots the counter-cumulative distribution function (in logs) observed in the simulated model compared to a Pareto and a lognormal distribution with coefficients estimated from the simulated model. As already mentioned, the model accounts only for part of the dispersion in firms' productivity, therefore in employment. However, it reproduces successfully the qualitative shape of the distribution's right tail. Although the literature has relied heavily on early results from Axtell (2001), according to which the firm-size distribution is well approximated by a Pareto distribution, recent evidence shows significant deviations from Pareto. When the goodness of fit is estimated through Maximum Likelihood, a more appropriate technique for these purposes, the right tail of the distribution falls in between a lognormal and a Pareto distribution (Kondo et al., 2023). The figure is also qualitatively in line with the one shown by Rossi-Hansberg and Wright (2007) for manufacturing establishments and firms in the US.

A further non-targeted moment is the ratio of production and non-supervisory employees to total employment based on BLS data. This moment is directly tied to the average firm size. In the model, production employment is 79% of overall employment, whereas it is 81% in the data.

Finally, as explained earlier, the model deviates from Gibrat Law, which has been the subject of extensive empirical investigation. In its simplest form, Gibrat's Law is tested by regressing firms' growth rates on their initial log size. Running this regression on one hundred thousand simulated observations, I obtain a coefficient of -0.027 . This modestly negative deviation is consistent with the empirical literature. Early influential studies on U.S. manufacturing include (Evans, 1987), which documents systematically negative size-growth elasticities (especially for smaller firms), and Hall (1987), which finds coefficients generally closer to zero (around -0.03 to -0.05) in a sample restricted to larger, publicly listed firms. Consistent with the model's predictions, more recent work using quantile regressions confirms that deviations from Gibrat Law are stronger among smaller firms and tend to diminish for larger firms (Distante et al., 2018).

Table 3
Results following a reduction and disappearance of highly innovative startups.

Moment	Model low skewness (−6.5 pp)	Model no skewness	Data
Δ Job reallocation incumb.	−0.35 pp	−2.25 pp	−2.3 pp
Δ Productivity growth	−0.15 pp	−0.65 pp	−0.8 pp
Δ Entry	−0.4 pp	−0.6 pp	−2.2 pp

5.2. The decline of fast growing startups

Decker et al. (2016b) document that one of the most salient changes at the turn of the new millennium is the disappearance of high-growth startups. This change is reflected in a reduction in the skewness of firm growth rates, especially among entrants — a moment targeted in the calibration. Specifically, the difference between the right tail (the 90th–50th percentile gap) and the left tail (50th–10th) declines by 6.5 percentage points between 1981–2002 and 2003–2011. I calibrate the parameter σ_α^E , which governs the dispersion of entrants' R&D productivity draws, to match this decline. Because the empirical change reflects the disappearance of highly innovative startups rather than a redistribution that preserves the mean, this calibration implies both lower dispersion and a lower average level of innovative ability among entrants — as χ_α is held constant —, reproducing the patterns documented by Decker et al. (2016b).

From a qualitative perspective, growth declines because the economy becomes less innovative. Average innovative ability among entrants falls as a result of both the disappearance of highly innovative startups and a potential reduction in overall entry and exit. Job reallocation among incumbents declines because highly innovative startups are typically small relative to the size that would allow them to grow at the same rate as the rest of the economy. Overall, because they operate far from their no-churning locus, they contribute disproportionately to both job reallocation and aggregate growth. Finally, entry is affected in opposite directions: it is negatively affected by a lower expected draw of innovative ability, and positively affected by the reduction in dispersion of the innovation draw.

Table 3 presents the results for the measures of turbulence and growth and compares them to their corresponding data moments. The first column reports the results obtained by changing σ_α^E to match the observed decline in skewness in the entrants' growth rate distribution. To assess the maximum potential impact of this mechanism, the second column reports the results from a counterfactual in which σ_α^E is reduced sufficiently to eliminate all skewness.

The results in the first column show that the observed reduction in skewness is insufficient to account for much of the decline in growth and turbulence measured in the data. While the reduction in σ_α^E contributes to the growth and dynamism slowdown, its quantitative role is limited.

The inability of the decline in entrants' R&D productivity skewness to explain much of the observed decline in growth and turbulence does not reflect an inadequacy of this channel. As shown by the results in the second column, a sufficiently large reduction in σ_α^E can generate a substantial decline in growth and job reallocation, though not in entry. Aside from entry — which does not respond strongly to changes in σ_α^E — the presence of fast-growing startups plays a quantitatively important role for growth and churning. The surprising result is that, although the disappearance of fast-growing startups is a salient feature of the data, the magnitude of the observed change is not large enough to generate appreciable macroeconomic consequences.

These results indicate that additional mechanisms must be at play. Channels operating primarily through changes in a small subset of firms — such as the right tail of entrant growth — are unlikely, at the magnitudes observed in the data, to generate sizeable changes in aggregate growth and turbulence. This motivates the exploration of explanations based on changes that affect a broader set of firms.

The literature has offered several explanations for the decline in business dynamism. A notable trend over the same period is the slowdown in labor force growth, which Karahan et al. (2024) estimate accounts for about half of the observed decline in entry rates — a result that this model can replicate, though without adding additional insights. As for the decline in churning, proposed mechanisms include reduced responsiveness to firm-specific shocks (Decker et al., 2020) and rising product market concentration associated with the emergence of “superstar firms” (Autor et al., 2020).

On the growth front, one set of explanations emphasizes changes in incumbent behavior associated with increased concentration (Ghazi, 2019; Olmstead-Rumsey, 2020; Ferraro et al., 2025). An alternative view holds that the slowdown reflects an endogenous response to the Great Recession (Anzoategui et al., 2019). Finally, another hypothesis is that the slowdown reflects the natural unwinding of a period of unusually fast growth. One possible interpretation is the end of the diffusion and adoption of information and communication technologies (ICTs), following the exhaustion of key complementary innovations (Fernald, 2015; Gordon, 2016). The next subsection considers a broader decline in innovative opportunities across firms, calibrated to match the observed slowdown in aggregate productivity growth, and examines whether the implied changes are also consistent with the observed decline in turbulence.

5.3. A broader decline in innovative opportunities

On top of the mechanism explored in the previous subsection, I next consider the possibility that the remaining productivity growth slowdown reflects a broader deterioration in innovative opportunities across firms. One possible interpretation of such a change is that it captures the exhaustion of the diffusion and adoption of ICTs, but this is only a motivating example rather than the

Table 4

Results following a reduction in R&D productivity of horizontal and vertical innovation for all firms. These results build upon the earlier analysis on the reduction in high growth startups, with the new parameters varied on top of the previous parameter change. The reduction in growth is the targeted moment.

Moment	Model low skewness (−6.5 pp)	Model low skewness and R&D productivity	Data
Δ Job reallocation incumb.	−0.35 pp	−2.6 pp	−2.3 pp
Δ Productivity growth	−0.15 pp	−0.8 pp	−0.8 pp
Δ Entry	−0.4 pp	−0.9 pp	−2.2 pp

Table 5

Asymmetric parameter changes.

Parameter varied	Change (log points)	Δ Reallocation (pp)	Δ Growth (pp)	Δ Entry (pp)
$\bar{\alpha}$	−0.1	−2.0	−0.6	−0.8
$\bar{\alpha}$	−0.3	−3.0	−1.0	−1.7
χ_Z	−0.1	−2.6	−0.7	−0.8
χ_Z	−0.3	−2.7	−0.85	−1.7
χ_α	−0.1	−2.4	−0.7	−0.8
χ_α	−0.3	−2.8	−0.9	−1.4

Notes: Starting from the symmetric benchmark, where $\bar{\alpha}$, χ_Z , and χ_α are each reduced by 0.2 log points, the table varies one parameter at a time while holding the other two fixed at their benchmark values.

specific hypothesis being tested. This subsection asks whether a generalized decline in innovative opportunities can also account for a reduction in turbulence that is quantitatively similar to that observed in the data. In contrast to the channel explored above, which operates through changes in the right tail of entrants' growth distribution, this subsection focuses on a broader mechanism affecting innovation by both incumbents and entrants.

In the model, three parameters regulate firms' innovativeness: $\bar{\alpha}$ governs the R&D productivity of incumbents, χ_Z affects the outcome of efforts in horizontal innovation, and χ_α determines the initial R&D productivity draw of entrants. The exercise reduces all three parameters by the same percentage, thereby imposing a symmetric decline in innovative opportunities across these margins. This symmetric specification is intended as an agnostic benchmark: rather than attributing the slowdown to a particular innovation channel, it captures a broad-based deterioration that affects incumbents and entrants alike.

While one could allow for asymmetric changes in these parameters to match individual moments more closely, doing so is not the objective of this section. Nevertheless, to assess whether the results depend on the exact symmetric benchmark, I consider a simple robustness exercise in the spirit of Pappa (2009), who varies calibrated parameters to evaluate the robustness of the model's implications. Specifically, starting from the baseline case in which $\bar{\alpha}$, χ_Z , and χ_α are each reduced by 0.2 log points, I vary each parameter one at a time while keeping the other two fixed at their benchmark values. For each parameter, I consider a smaller reduction of 0.1 log points and a larger reduction of 0.3 log points. This yields a set of asymmetric perturbations around the symmetric benchmark, which allows me to assess whether the conclusions of this subsection depend materially on the exact composition of the decline in innovative opportunities across incumbents, entrants, and horizontal innovation.

The reason a reduction in R&D productivity reduces turbulence is straightforward to understand by considering the extreme case of no firm growth. If firms do not increase their technological knowledge, they do not expand in size. Moreover, if aggregate productivity growth is zero, no firm shrinks either, since shrinkage in the model arises from firms losing technological ground relative to the rest of the economy.

Turning to magnitudes, Table 4 reports the quantitative results. The parameter changes slightly overexplain the decline in job reallocation among incumbents, indicating that this margin responds strongly to broad-based changes in R&D productivity.

The effect on entry is quantitatively smaller, though still meaningful. According to BDS data, employment growth declines by 0.9 percentage points over this period, which mechanically implies an equal reduction in the entry rate. The parameter changes account for roughly two-thirds of the remaining 1.3 percentage point decline in entry.

Taken together, these results indicate that the joint decline in productivity growth and turbulence observed in the data can plausibly be reconciled with a generalized slowdown in innovative activity. The asymmetric parameter changes in Table 5 support the same conclusion.

6. Summary and conclusions

This paper develops a unified framework to study aggregate productivity growth and turbulence. Firms are monopolistically competitive, invest in in-house R&D, and face idiosyncratic innovation shocks. Competition for market share implies diminishing returns in relative terms, generating endogenous churning as firms adjust their size, while increasing returns in absolute terms sustain positive long-run growth. Entry and exit introduce life-cycle dynamics that further shape turbulence and aggregate productivity by continuously replacing goods in the market.

The paper 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). I add endogenous growth to the former, and turbulence to the latter.

In a quantitative analysis, I examine the impact of the disappearance of highly innovative startups, as documented by Decker et al. (2016b) using U.S. data. Although this mechanism has the potential to reduce growth and turbulence substantially, matching the observed reduction in the skewness of startups' growth rates does not generate quantitatively large effects. The calibrated model indicates that the thinning of the right tail of the distribution of innovation ability among entrants can account for 15% of the productivity growth slowdown that occurred in the mid-2000s, and for a similar share of the decline in entry and job reallocation among incumbents.

This result is informative, as it implies that, despite being a salient microeconomic fact, the disappearance of fast-growing startups is not quantitatively large enough to generate appreciable macroeconomic changes. By contrast, mechanisms that operate through a broader decline in innovative opportunities across firms – such as a generalized reduction in R&D productivity – appear more capable of jointly accounting for the observed evolution of growth, job reallocation, and entry. A contribution of the paper is therefore to discipline, within a unified framework, the quantitative importance of fast-growing startups for aggregate outcomes. More broadly, the analysis also suggests that mechanisms operating through a wider decline in innovative opportunities are better suited to generate sizeable macroeconomic effects.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Stationary model

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

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

The production function (9) is:

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

where the first-order condition for production labor, obtained from the corresponding level equation by multiplying and dividing by $LN^{\frac{1}{\epsilon-1}} N^{\frac{\epsilon}{\epsilon-1}} Z^{\theta(\epsilon-1)}$, is

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

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

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

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

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

The stationary Bellman equation is:

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

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

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

The entry condition (18) is:

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

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

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

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

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

output (21) is:

$$\hat{y}_t = \hat{c}_t + f_E \frac{N_{E_t}}{N_t} n_t; \tag{A.10}$$

Appendix B. Steady state algorithm

Construct a grid for the state z (relative knowledge) and the shock α (productivity of R&D) by choosing respectively 230 and 120 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 Eq. (A.5) via policy function iteration for the R&D labor of firms at each combination of grid points of the two state variables, subject to the constraint (A.6).

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. If the firm value is below 0, set production and R&D labor to 0, as the firm exits the market.

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:

- Interpolate the relevant variables on a grid with 2000 points for $z_{i,t}$ and 300 points for $\alpha_{i,t}$.
- 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 Eq. (16) and $\alpha_{i,t}$ according to Eq. (17).
- Find the entry rate as the sum of exit rate and population growth rate (the condition required to ensure stationarity in the number of firms, essentially imposing steady state) from Eq. (A.9).
- Compute the new mass of firms as the weighted average of the mass of incumbents and the mass of entrants, using the entry rate as the weight.
- Iterate until the mass of firms in every grid point is close enough from what it was in the previous iteration.

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

Appendix C. Proof of Proposition 2

With these two assumptions, (31) simplifies considerably. First, the expectation operator becomes unnecessary as the problem is now one of perfect foresight. Then the last term in the square bracket drops out. As it is easier to work with stationarized variables, specifically, I define for any variable X_t , $\hat{X}_t = X_t / Z_t^\theta N_t^{\frac{1}{c-1}}$ (Appendix A provides all relevant equations with only stationary variables). Iterating forward that equation then leads to:

$$l_{Z_{i,t}} = \left[\frac{B_{t+1} \alpha_{i,t} z_{i,t}^\mu \zeta}{\hat{w}_t} \frac{\partial \hat{\pi}_{i,t+1}}{\partial z_{i,t+1}} \right]^{\frac{1}{1-\zeta}} \tag{C.1}$$

where

$$B_{t+1} = \frac{(1 + g_{t+1}^Z)^{\theta-1} (1 + \lambda)^{\frac{1}{c-1}} \left(\frac{n_{t+1}}{n_t} \right)^{\frac{1}{c-1}}}{1 + r_{t+1}}. \tag{C.2}$$

Because Eq. (C.1) depends simultaneously on z_t and z_{t+1} , I iterate the z_{t+1} term backwards to write the equation only as a function of the current level of technological knowledge. At this point, the expression for $l_{z_{i,t}}$ can be substituted into the expression for the firm's productivity growth rate, given in Eq. (30), and then take the derivative with respect to $z_{i,t}$. As I will show, that derivative is negative when the two requirements pointed out in Proposition 2 are satisfied.

Optimal knowledge growth is:

$$g_{i,t+1}^Z = G_0 z_{i,t}^{\frac{\mu}{1-\zeta}-1} \frac{[\theta(\epsilon-1)-1] \frac{\zeta}{1-\zeta}}{z_{i,t+1}} \tag{C.3}$$

with G_0 being a positive constant equal to:

$$G_0 = \alpha_{i,t}^{1+\frac{\zeta}{1-\zeta}} \left[B_{t+1} \zeta \left(\frac{\epsilon-1}{\epsilon} \right)^\epsilon \hat{w}_{t+1}^{1-\epsilon} \hat{y}_{t+1} \right]^{\frac{\zeta}{1-\zeta}} \tag{C.4}$$

Iterating $z_{i,t+1}$ backwards, it can be rewritten as

$$g_{i,t+1}^Z = G_0 z_{i,t}^{\frac{\mu}{1-\zeta}-1+\theta(\epsilon-1)-1} \frac{\zeta}{1-\zeta} \left(1 + g_{i,t+1}^Z \right)^{[\theta(\epsilon-1)-1] \frac{\zeta}{1-\zeta}} \tag{C.5}$$

The term $1 + g_{i,t+1}^Z$ appears because of the dependence of optimal R&D on relative knowledge both at time t and at time $t + 1$. In a continuous time version of this model, $t \approx t + 1$, thus eliminating that term. In that case, the first requirement for stationarity in Proposition 2 would be the only one needed. Instead, because time is discrete, implicit differentiation is needed to find the derivative. The result is:

$$\frac{\partial g_{i,t+1}^Z}{\partial z_{i,t}} = \frac{G_0 z_{i,t}^{(\cdot)-1} \left(1 + g_{i,t+1}^Z \right)^{(\cdot)} [\zeta \theta (\epsilon - 1) - 1 + \mu]}{1 - G [\theta (\epsilon - 1) - 1]} \tag{C.6}$$

where G is the positive constant

$$G = G_0 \frac{\zeta}{1-\zeta} z_{i,t}^{(\cdot)} \left(1 + g_{i,t+1}^Z \right)^{(\cdot)-1} \tag{C.7}$$

Next, I prove condition (ii). Since knowledge growth is strictly decreasing, and growth of average knowledge is strictly positive and finite, a sufficient condition for this proof is that knowledge growth tends to infinity when relative knowledge tends to zero, and it tends to zero when relative knowledge tends to infinity. I proceed by proving each part in sequence.

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

$$\lim_{z_{i,t} \rightarrow \infty} z_{i,t+1} = \frac{z_{i,t} + \alpha_{i,t} z_{i,t}^{\frac{\mu}{1-\zeta}} l_{z_{i,t}}^{\zeta}}{1 + g_{t+1}^Z} = \infty, \tag{C.8}$$

for any value of $l_{z_{i,t}}$, given its non-negativity constraint. At this stage, we can find the limit of $g_{i,t+1}^Z$ when both $z_{i,t}$ and $z_{i,t+1}$ tend to infinity. From Eq. (C.3), it is immediate to see that the growth rate of knowledge equals 0 when the sum of the two exponents is negative. The sum of the two exponents is the first requirement for stationarity in Proposition 2.

For the second part of condition (ii), we need to prove that knowledge growth tends to infinity when the relative knowledge level tends to 0. Whenever R&D investment is positive, it is clear that knowledge growth tends to infinity because it is an increasing function of R&D investment and relative knowledge is the only element that features at the denominator because $\mu < 1$ by assumption. Instead, suppose that the R&D investment is 0. In this case, the logic follows the one used to prove the first part of condition (ii). That is, if $z_{i,t}$ tends to 0, so does $z_{i,t+1}$. Hence whenever the parameters satisfy the requirement for stationarity, from Eq. (C.3) we see that growth must tend to infinity.

Appendix D. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.eurocorev.2026.105331>.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., Kerr, W., 2018. Innovation, reallocation, and growth. *Am. Econ. Rev.* 108, 3450–3491. <http://dx.doi.org/10.1257/aer.20130470>.

Acemoglu, D., Cao, D., 2015. Innovation by entrants and incumbents. *J. Econom. Theory* 157, 255–294. <http://dx.doi.org/10.1016/j.jet.2015.01.001>.

Acemoglu, D., Ventura, J., 2002. The world income distribution. *Q. J. Econ.* 117, 659–694.

Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. *Econometrica* 60, 323–351. <http://dx.doi.org/10.2307/2951599>.

Akcigit, U., Kerr, W.R., 2018. Growth through heterogeneous innovations. *J. Political Econ.* 126, 1374–1443. <http://dx.doi.org/10.1086/697901>.

Anzoategui, D., Comin, D., Gertler, M., Martinez, J., 2019. Endogenous technology adoption and R&D as sources of business cycle persistence. *Am. Econ. J.: Macroecon.* 11, 67–110. <http://dx.doi.org/10.1257/mac.20170269>.

Audretsch, D., Klomp, L., Santarelli, E., Thurik, A., 2004. Gibrat's law: Are the services different? *Rev. Ind. Organ.* 24, 301–324. <http://dx.doi.org/10.1023/B:REIO.0000038273.50622.ec>.

Autor, D., Dorn, D., Katz, L.F., Patterson, C., Van Reenen, J., 2020. The fall of the labor share and the rise of superstar firms. *Q. J. Econ.* 135, 645–709. <http://dx.doi.org/10.1093/qje/qjaa004>.

- Axtell, R.L., 2001. Zipf distribution of U.S. Firm sizes. *Science* 293, 1818–1820. <http://dx.doi.org/10.1126/science.1062081>.
- Bilbiie, F.O., Ghironi, F., Melitz, M.J., 2012. Endogenous entry, product variety, and business cycles. *J. Political Econ.* 120, 304–345. <http://dx.doi.org/10.1086/665825>.
- Bloom, N., Schankerman, M., Van Reenen, J., 2013. Identifying technology spillovers and product market rivalry. *Econometrica* 81, 1347–1393. <http://dx.doi.org/10.3982/ECTA9466>.
- Bottazzi, G., Cefis, E., Dosi, G., Secchi, A., 2007. Invariances and diversities in the patterns of industrial evolution: Some evidence from Italian manufacturing industries. *Small Bus. Econ.* 29, 137–159. <http://dx.doi.org/10.1007/s11187-006-0014-y>.
- Brown, C., Haltiwanger, J., Lane, J., 2008. *Economic Turbulence: Is a Volatile Economy Good for America?* University of Chicago Press.
- Byrne, D.M., Fernald, J.G., Reinsdorf, M.B., 2016. Does the United States have a productivity slowdown or a measurement problem? *Brook. Pap. Econ. Act.* 2016, 109–182. <http://dx.doi.org/10.1353/eca.2016.0014>.
- De Ridder, M., 2024. Market power and innovation in the intangible economy. *Am. Econ. Rev.* 114, 199–251. <http://dx.doi.org/10.1257/aer.20201079>.
- De Ridder, M., Grassi, B., Morzenti, G., 2022. *The Hitchhiker's guide to markup estimation*.
- Decker, R.A., Haltiwanger, J., Jarmin, R.S., Miranda, J., 2016a. Declining business dynamism: What we know and the way forward. *Am. Econ. Rev.* 106, 203–207. <http://dx.doi.org/10.1257/aer.p20161050>.
- Decker, R.A., Haltiwanger, J., Jarmin, R.S., Miranda, J., 2016b. Where has all the skewness gone? The decline in high-growth (young) firms in the U.S. *Eur. Econ. Rev.* 86, 4–23. <http://dx.doi.org/10.1016/j.eurocorev.2015.12.013>.
- Decker, R.A., Haltiwanger, J., Jarmin, R.S., Miranda, J., 2020. Changing business dynamism and productivity: Shocks versus responsiveness. *Am. Econ. Rev.* 110, 3952–3990. <http://dx.doi.org/10.1257/aer.20190680>.
- Dinopoulos, E., Thompson, P., 1998. Schumpeterian growth without scale effects. *J. Econ. Growth* 3, 313–335. <http://dx.doi.org/10.1023/A:1009711822294>.
- Distante, R., Petrella, I., Santoro, E., 2018. Gibrat's law and quantile regressions: An application to firm growth. *Econom. Lett.* 164, 5–9. <http://dx.doi.org/10.1016/j.econlet.2017.12.028>.
- Dosi, G., 1988. Sources, procedures, and microeconomic effects of innovation. *J. Econ. Lit.* 26, 1120–1171.
- Ericson, R., Pakes, A., 1995. Markov-perfect industry dynamics: A framework for empirical work. *Rev. Econ. Stud.* 62, 53–82. <http://dx.doi.org/10.2307/2297841>.
- Evans, D.S., 1987. The relationship between firm growth, size, and age: Estimates for 100 manufacturing industries. *J. Ind. Econ.* 35, 567–581. <http://dx.doi.org/10.2307/2098588>.
- Fernald, J.G., 2015. Productivity and potential output before, during, and after the great recession. *NBER Macroecon. Annu.* 29 (1), 1–51.
- Fernald, J., 2018. Is slow productivity and output growth in advanced economies the new normal? *Int. Prod. Monit.* 35, 138–148.
- Ferraro, D., Ghazi, S., Peretto, P.F., 2025. Fixed costs, markups, and productivity growth.
- Foster, L., Haltiwanger, J., Krizan, C., 2001. Aggregate productivity growth: Lessons from microeconomic evidence. In: *New Developments in Productivity Analysis*. University of Chicago Press, Chicago, pp. 303–363. <http://dx.doi.org/10.3386/w6803>.
- Foster, L., Haltiwanger, J., Syverson, C., 2008. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *Am. Econ. Rev.* 98, 394–425. <http://dx.doi.org/10.1257/aer.98.1.394>.
- Garcia-Macia, D., Hsieh, C.T., Klenow, P.J., 2019. How destructive is innovation? *Econometrica* 87, 1507–1541. <http://dx.doi.org/10.3982/ECTA14930>.
- Ghazi, S., 2019. *Large firms and long-run growth*.
- Gomme, P., Ravikumar, B., Rupert, P., 2011. The return to capital and the business cycle. *Rev. Econ. Dyn.* 14, 262–278. <http://dx.doi.org/10.1016/j.red.2010.11.004>.
- Gordon, R.J., 2016. *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*. Princeton University Press.
- Ha, J., Howitt, P., 2007. Accounting for trends in productivity and R&D: A schumpeterian critique of semi-endogenous growth theory. *J. Money Credit. Bank.* 39, 733–774. <http://dx.doi.org/10.1111/j.1538-4616.2007.00045.x>.
- Hall, B.H., 1987. The relationship between firm size and firm growth in the US manufacturing sector. *J. Ind. Econ.* 35, 583–606. <http://dx.doi.org/10.2307/2098589>.
- Hall, B.H., Lerner, J., 2010. Chapter 14 - the financing of R&D and innovation. In: Hall, B.H., Rosenberg, N. (Eds.), *Handbook of the Economics of Innovation*. vol. 1, North-Holland, pp. 609–639. [http://dx.doi.org/10.1016/S0169-7218\(10\)01014-2](http://dx.doi.org/10.1016/S0169-7218(10)01014-2).
- Hopenhayn, H.A., 1992. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica* 60, 1127–1150. <http://dx.doi.org/10.2307/2951541>.
- Hopenhayn, H.A., 2014. Firms, misallocation, and aggregate productivity: A review. *Annu. Rev. Econ.* 6, 735–770. <http://dx.doi.org/10.1146/annureconomics-082912-110223>.
- Karahan, F., Pugsley, B., Şahin, A., 2024. Demographic origins of the start-up deficit. *Am. Econ. Rev.* 114, 1986–2023. <http://dx.doi.org/10.1257/aer.20210362>.
- Klette, T.J., Kortum, S., 2004. Innovating firms and aggregate innovation. *J. Political Econ.* 112, 986–1018. <http://dx.doi.org/10.1086/422563>.
- Kondo, I.O., Lewis, L.T., Stella, A., 2023. Heavy tailed but not Zipf: Firm and establishment size in the United States. *J. Appl. Econometrics* 38, 767–785. <http://dx.doi.org/10.1002/jae.2976>.
- Laincz, C.A., 2009. R&D subsidies in a model of growth with dynamic market structure. *J. Evol. Econ.* 19, 643–673. <http://dx.doi.org/10.1007/s00191-008-0114-8>.
- Laincz, C.A., Peretto, P.F., 2006. Scale effects in endogenous growth theory: An error of aggregation not specification. *J. Econ. Growth* 11, 263–288. <http://dx.doi.org/10.1007/s10887-006-9004-9>.
- Lee, Y., Mukoyama, T., 2015. Entry and exit of manufacturing plants over the business cycle. *Eur. Econ. Rev.* 77, 20–27. <http://dx.doi.org/10.1016/j.eurocorev.2015.03.011>.
- Lentz, R., Mortensen, D.T., 2008. An empirical model of growth through product innovation. *Econometrica* 76, 1317–1373. <http://dx.doi.org/10.3982/ECTA5997>.
- Lentz, R., Mortensen, D.T., 2016. Optimal growth through product innovation. *Rev. Econ. Dyn.* 19, 4–19. <http://dx.doi.org/10.1016/j.red.2015.12.002>.
- Luttmer, E.G.J., 2007. Selection, growth, and the size distribution of firms. *Q. J. Econ.* 122, 1103–1144. <http://dx.doi.org/10.1162/qjec.122.3.1103>.
- Madsen, J.B., 2008. Semi-endogenous versus schumpeterian growth models: Testing the knowledge production function using international data. *J. Econ. Growth* 13, 1–26. <http://dx.doi.org/10.1007/s10887-007-9024-0>.
- Mand, M., 2019. On the cyclicality of R&D activities. *J. Macroecon.* 59, 38–58. <http://dx.doi.org/10.1016/j.jmacro.2018.10.008>.
- Massari, F., 2025. Revisiting productivity growth accounting decompositions. *Res. Econ.* 79, 101055. <http://dx.doi.org/10.1016/j.rie.2025.101055>.
- Massari, F., Peretto, P.F., 2026. *Super-robust endogenous growth: Theory and empirical insights*. *J. Econ. Behav. Organ.* 245, 107549.
- Melitz, M.J., Polanec, S., 2015. Dynamic Olley–Pakes productivity decomposition with entry and exit. *Rand J. Econ.* 46, 362–375. <http://dx.doi.org/10.1111/1756-2171.12088>.
- Naudé, W., 2022. From the entrepreneurial to the ossified economy. *Camb. J. Econ.* 46, 105–131.
- Olmstead-Rumsey, J., 2020. Market concentration and the productivity slowdown.
- Pappa, E., 2009. The effects of fiscal shocks on employment and the real wage. *Internat. Econom. Rev.* 50 (1), 217–244.
- Peretto, P.F., 1998. Technological change, market rivalry, and the evolution of the capitalist engine of growth. *J. Econ. Growth* 3, 53–80. <http://dx.doi.org/10.1023/A:1009722031825>.
- Peretto, P.F., 1999. Cost reduction, entry, and the interdependence of market structure and economic growth. *J. Monet. Econ.* 43, 173–195. [http://dx.doi.org/10.1016/S0304-3932\(98\)00040-3](http://dx.doi.org/10.1016/S0304-3932(98)00040-3).
- Peretto, P.F., 2018. Robust endogenous growth. *Eur. Econ. Rev.* 108, 49–77. <http://dx.doi.org/10.1016/j.eurocorev.2018.06.007>.
- Peretto, P.F., 2026. *Knowledge, Industry and the Endogenous Growth of Nations*. MIT Press, forthcoming.
- Peretto, P.F., Connolly, M., 2007. The manhattan metaphor. *J. Econ. Growth* 12, 329–350. <http://dx.doi.org/10.1007/s10887-007-9023-1>.

- Peretto, P.F., Smulders, S., 2002. Technological distance, growth and scale effects. *Econ. J.* 112, 603–624. <http://dx.doi.org/10.1111/1468-0297.00732>.
- Peters, M., 2020. Heterogeneous markups, growth, and endogenous misallocation. *Econometrica* 88, 2037–2073. <http://dx.doi.org/10.3982/ECTA15565>.
- Restuccia, D., Rogerson, R., 2017. The causes and costs of misallocation. *J. Econ. Perspect.* 31, 151–174. <http://dx.doi.org/10.1257/jep.31.3.151>.
- Rossi-Hansberg, E., Wright, M.L.J., 2007. Establishment size dynamics in the aggregate economy. *Am. Econ. Rev.* 97, 1639–1666. <http://dx.doi.org/10.1257/aer.97.5.1639>.
- Sutton, J., 1997. Gibrat's legacy. *J. Econ. Lit.* 35, 40–59.
- Syverson, C., 2011. What determines productivity? *J. Econ. Lit.* 49, 326–365. <http://dx.doi.org/10.1257/jel.49.2.326>.
- Syverson, C., 2017. Challenges to mismeasurement explanations for the US productivity slowdown. *J. Econ. Perspect.* 31, 165–186. <http://dx.doi.org/10.1257/jep.31.2.165>.
- Thompson, P., 2001. The microeconomics of an R&D-based model of endogenous growth. *J. Econ. Growth* 6, 263–283. <http://dx.doi.org/10.1023/A:1012761811439>.
- Young, A., 1998. Growth without scale effects. *J. Political Econ.* 106, 41–63. <http://dx.doi.org/10.1086/250002>.