Australian National University
Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Heterogeneous Firms, Growth and the Long Shadows of Business Cycles

CAMA Working Paper 93/2026
November 2026

Cristiana Benedetti-Fasil
RC European Commission*

Giammario Impullitti
University of Nottingham

Omar Licandro
University of Leicester

Petr Sedláček
University of New South Wales

Adam Hal Spencer
University of Bonn
Centre for Applied Macroeconomic Analysis, ANU

Abstract

Do firms differ in their innovation response to business cycles? What economic forces shape this heterogeneity? And does it matter for aggregate propagation? We document substantial heterogeneity in firm-level R&D cyclicality, develop a theory to explain this fact, and validate its key mechanism in the data. In a tractable model of business cycles and innovation-driven growth, endogenous selection makes R&D particularly responsive to aggregate conditions for firms close to the exit margin. Consistent with this mechanism, firms closer to exit display more procyclical R&D in the data. Quantitatively, selection and endogenous growth interact powerfully: selection generates heterogeneous R&D cyclicality, while endogenous growth translates these heterogeneous responses into persistent aggregate effects. Ignoring heterogeneous R&D cyclicality misses roughly half of the long-run output cost of a severe recession.

Keywords

growth, business cycles, innovation, heterogeneous firms, selection

JEL Classification

E32, O31, O40, L1

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

The Centre for Applied Macroeconomic Analysis in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

Heterogeneous Firms, Growth and the Long Shadows of Business Cycles*

Cristiana Benedetti-Fasil Giammario Impullitti Omar Licandro
Petr Sedláček Adam Hal Spencer

October 27, 2026

Abstract

Do firms differ in their innovation response to business cycles? What economic forces shape this heterogeneity? And does it matter for aggregate propagation? We document substantial heterogeneity in firm-level R&D cyclicality, develop a theory to explain this fact, and validate its key mechanism in the data. In a tractable model of business cycles and innovation-driven growth, endogenous selection makes R&D particularly responsive to aggregate conditions for firms close to the exit margin. Consistent with this mechanism, firms closer to exit display more procyclical R&D in the data. Quantitatively, selection and endogenous growth interact powerfully: selection generates heterogeneous R&D cyclicality, while endogenous growth translates these heterogeneous responses into persistent aggregate effects. Ignoring heterogeneous R&D cyclicality misses roughly half of the long-run output cost of a severe recession.

Keywords: Growth, business cycles, innovation, heterogeneous firms, selection

JEL Codes: E32, O31, O40, L1

*Our wonderful friend and colleague Cristiana Benedetti-Fasil is no longer with us. A previous version of the paper was titled "Heterogeneous Firms, R&D Policies and the Long Shadows of Business Cycles".

Email addresses: Impullitti (University of Nottingham) giammario.impullitti@nottingham.ac.uk, Licandro (University of Leicester) omar.licandro@leicester.ac.uk, Sedláček (University of New South Wales) p.sedlacek@unsw.edu.au, Spencer (University of Bonn & ANU-CAMA) spencer@uni-bonn.de.

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project C05) and Germany's Excellence Strategy – EXC 2126/2 – 390838866 is gratefully acknowledged. We are grateful for comments received during talks at Carnegie Mellon University, New York Fed, Penn State University, Chicago Fed, Boston Fed, Queen's University, University of Toronto, University of Western Ontario, Cornell University, UCL, Bank of England, Banque de France, University of Edinburgh, University of St. Andrews, Central Bank of Ireland, Philadelphia Fed, Cleveland Fed, SED Athens 2026, RBA, University of Sydney, Richmond Fed, St Louis Fed, Ohio State, Deakin University, QUT, De Nederlandsche Bank, Bank of Finland, Helsinki GSE, Berlin Macro Seminar, Econometric Society European Winter Meeting 2025, Midwest Macro Meeting Fall 2025 & 2025 Workshop of the Spanish Macro Network.

1 Introduction

Recent events, from the Great Recession to the Covid-19 pandemic, have renewed concerns that severe cyclical downturns may leave persistent economic scars by affecting the economy's long-run growth path. Indeed, many countries have failed to return to their pre-crisis output trajectories following the Great Recession (Ball, 2014; Blanchard et al., 2015).[1] One natural channel through which temporary shocks can have long-lived effects is innovation. Unlike other investment margins, a temporary disruption to innovation can alter the subsequent path of productivity itself, allowing short-lived shocks to leave persistent scars. Innovation is inherently heterogeneous across firms: firms differ sharply in their innovative effort, and aggregate productivity growth is the outcome of these cross-sectional innovation decisions. This heterogeneity is a central feature of modern theories of long-run growth (Klette and Kortum, 2004; Akcigit and Kerr, 2018; Acemoglu et al., 2018).

Innovation is also procyclical in the data.[2] Downturns therefore reduce not only current economic activity but also investment in future productivity, providing a natural mechanism through which temporary shocks can cast long shadows. Do firms differ in how strongly their innovation responds to the cycle? What economic forces shape this heterogeneity? And does it matter for aggregate propagation? We address these questions in three steps. First, we document substantial heterogeneity in the cyclicality of firm-level R&D. Second, we develop a theory that explains why firms differ in their innovation responses to aggregate conditions and validate its key mechanism in the data. Finally, we use a calibrated version of the model to assess how heterogeneous innovation cyclicality shapes the persistent output consequences of severe downturns.

Compustat data reveal substantial dispersion in firm-level R&D cyclicality, with R&D procyclical on average. We show theoretically that this pattern does not follow mechanically from firm heterogeneity. Firms can differ substantially in productivity, size, profits, and steady-state R&D while still having the same innovation elasticity to an aggregate shock. Endogenous selection breaks this neutrality and predicts that firms closer to the exit margin have more cyclical R&D. We test this prediction directly and find that firms closer to exit indeed display significantly more procyclical R&D.

Quantitatively, selection and endogenous growth interact powerfully. Selection alone adds relatively little to long-run propagation when productivity growth is exogenous, but becomes quantitatively important when firms' innovation decisions determine future productivity. Ignoring the heterogeneous R&D cyclicality generated by selection misses

[1] The observation that recoveries may not always be robust enough to return GDP to its pre-crisis trend extends beyond this particular crisis. Empirical evidence abundantly shows that GDP fluctuations are persistent, with their effects lingering for years after the initial shock. See Cerra et al. (2023) for a recent survey.

[2] Empirical evidence on the procyclicality of aggregate R&D can be found in Griliches (1990); Comin and Gertler (2006); Barlevy (2007); Ouyang (2011), among others.

roughly half of the long-run output scar of a severe recession. More broadly, the aggregate consequences of a recession depend not only on the fall in innovation, but also on the cross-sectional pattern of firms' innovation responses.

Our framework combines a standard heterogeneous-firm environment (e.g. Hopenhayn, 1992) with innovation-driven growth. A continuum of heterogeneous firms produces differentiated goods under CES demand. Incumbents and prospective entrants invest in R&D to improve firm productivity, and their innovation decisions jointly determine aggregate productivity growth. Firms also face stochastic fixed operating costs and endogenously exit when the value of continuing falls below the realized cost. Aggregate productivity shocks affect firms' cashflows and hence their incentives to innovate and remain in the market. The environment is deliberately parsimonious: firms face the same aggregate conditions, and we introduce no firm-specific cyclical shocks, asymmetric financing frictions, or heterogeneous adjustment costs to generate differences in their responses.

The mechanism can be seen transparently in a simple two-period version of the model. Under CES demand, an aggregate shock generates the same first-order percentage change in gross cashflows for all firms, regardless of their productivity. Absent endogenous exit, the value gain from successful innovation therefore moves proportionally with aggregate conditions, implying a common first-order innovation elasticity across firms. Selection breaks this neutrality because successful innovation not only raises future profits but also increases the probability of survival. As a result, the value gain from innovation becomes more sensitive to aggregate conditions, amplifying firms' R&D responses. This channel is strongest for firms near the exit margin, for whom innovation can make the difference between survival and exit.

We take this prediction directly to the data. Using Compustat firm-level R&D together with industry-level output growth, we estimate how R&D cyclicality varies with firms' proximity to exit. We use two complementary measures. The first identifies firms in the year immediately preceding exit through bankruptcy or liquidation. The second uses distance to default to measure ex ante exposure to the exit margin. Both deliver the same pattern: firms closer to exit display significantly more procyclical R&D. The relationship is robust to time effects and a rich set of balance-sheet controls, suggesting that it is not simply capturing differential exposure to financing conditions.

We bring the full model to the data, calibrating its stochastic dynamics to discipline the selection–innovation mechanism. In particular, we target the relationship between R&D cyclicality and exit risk among firms accounting for a meaningful share of aggregate sales, together with average R&D cyclicality and standard moments of firm dynamics, innovation, and growth. Crucially, the calibrated model also reproduces the key empirical prediction of our selection mechanism without directly targeting it: firms closer to the exit margin display more procyclical R&D. The model generates about two thirds of the

corresponding empirical gradient, providing an important untargeted validation of the mechanism.

Armed with the calibrated model, we first confirm that the key prediction of the simple two-period model carries over to the full dynamic economy: firms closer to the exit margin exhibit more cyclical R&D. We then quantify how the interaction between endogenous selection and innovation-driven growth shapes the long-run effects of aggregate shocks. To disentangle these forces, we compare four economies according to whether productivity growth is endogenous or exogenous and whether endogenous selection is present or absent.

When growth is exogenous, shock propagation is substantially smaller than in the baseline, and introducing selection has only a modest effect despite sizeable movements in firm demographics. When innovation is endogenous but selection is removed, temporary shocks affect future productivity, as in standard endogenous-growth business-cycle models, but the resulting long-run scar is also substantially smaller than in the baseline. Importantly, this economy is calibrated to reproduce approximately the same average R&D cyclicality as the baseline, while eliminating the heterogeneity in firm-level responses generated by selection. The key quantitative finding is that the *interaction* between selection and endogenous growth generates substantially greater persistence than either mechanism in isolation. At a 50-year horizon, an endogenous-growth economy without selection accounts for only about half of the baseline output scar. This highlights that the long-run consequences of a recession depend not only on the average responsiveness of innovation, but also on the cross-sectional distribution of firms' innovation responses.

Literature review. Our paper relates first to the literature linking endogenous growth and business-cycle fluctuations. A large body of work shows that cyclical movements in innovation and technical change can propagate temporary shocks by altering the subsequent path of productivity (e.g. Fatas, 2000; Comin and Gertler, 2006; Barlevy, 2007; Queralto, 2020; Benigno and Fornaro, 2018; Anzoategui et al., 2019; Fornaro and Wolf, 2023; Bertolotti and Lanteri, 2024). Barlevy (2007) is our closest conceptual benchmark. Like us, he studies how standard aggregate productivity fluctuations affect innovation and, through it, long-run growth. In his framework, however, firms are symmetric, so equilibrium innovation is identical across firms both in its level and in its response to aggregate conditions. We contribute by characterizing the cross-sectional distribution of innovation responses in a heterogeneous-firm environment. We show that heterogeneity in productivity, size, profits, and steady-state R&D is not by itself sufficient to generate heterogeneous innovation elasticities. Our key theoretical result is that endogenous selection is the driver of this heterogeneity.

Our paper is also related to work in which financial frictions shape the cyclical behavior of productivity-enhancing investment. Aghion et al. (2010) show theoretically that credit constraints can make productivity-enhancing investment more procyclical. The

closest empirical precedent is Aghion et al. (2012), who use French firm-level data to show that R&D is more procyclical among firms facing tighter credit constraints. Their evidence therefore identifies heterogeneity in R&D cyclicality along a specific observable dimension—financial constraints—whereas we document the broader cross-sectional distribution of firm-level R&D elasticities and, more importantly, a new relationship between R&D cyclicality and proximity to the exit margin.

Related work shows that financial frictions affecting innovation can generate persistent output and productivity losses following financial crises (e.g. Queralto, 2020; Bianchi et al., 2019; Schmitz, 2021). Within heterogeneous-firm endogenous-growth models, Schmitz (2021) generates heterogeneous innovation responses because small firms and entrants are differentially exposed to financing constraints, while Ates and Saffie (2021) show that tighter financial conditions reduce entry but select toward firms with higher innovation efficiency, potentially mitigating the aggregate effects of financial shocks. Our mechanism is different. We consider a common aggregate productivity shock and abstract from borrowing constraints. Heterogeneous innovation responses arise from the interaction between endogenous selection and firms' incentives to innovate, rather than from heterogeneous exposure to financial shocks.

Finally, our paper relates to the broader literature on heterogeneous firm responses and business-cycle propagation. An established literature documents substantial cross-sectional differences in firms' cyclical responses (e.g. Gertler and Gilchrist, 1994; Fort et al., 2013; Crouzet and Mehrotra, 2020). This literature focuses on heterogeneity in outcomes such as sales, employment, and physical investment. Related work provides theoretical interpretation of this evidence based on heterogeneous access to credit and capital adjustment costs. For example, Ottonello and Winberry (2020) find that firms further from default respond more strongly to monetary-policy shocks through heterogeneity in the sensitivity of borrowing costs, while Cooper and Haltiwanger (2006) and Winberry (2021) emphasize nonlinear adjustment in physical capital. We contribute to this literature by studying heterogeneous innovation responses, identifying a selection-based mechanism that generates this heterogeneity, and quantifying its implications for the long-run path of productivity.

A related strand of this literature emphasizes propagation through firm entry, exit, and composition. A particularly relevant contribution is Clementi and Palazzo (2016), who show that aggregate shocks generate persistence by changing the mass and composition of firms. A positive shock generates a larger cohort of relatively low-productivity young firms, whose subsequent productivity growth reflects mean reversion in an exogenous productivity process. Their propagation mechanism therefore does not rely on selection affecting firms' investment decisions. In our model, by contrast, selection changes firms' incentives to innovate and thereby endogenously shapes the subsequent evolution of productivity. Our exogenous-growth counterfactual makes this distinction explicit: selection

and firm demographics alone generate only modest additional long-run propagation in our calibrated economy, whereas their interaction with endogenous innovation substantially magnifies the long-run effects of temporary shocks.

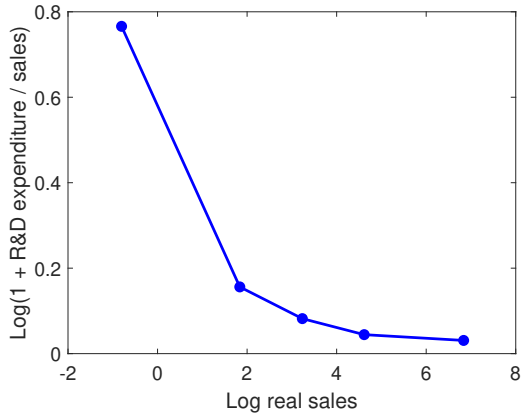
Taken together, our analysis makes four contributions. We document substantial heterogeneity in firm-level R&D cyclicality, establish theoretically that endogenous selection generates heterogeneous innovation responses to common aggregate shocks, and directly test this mechanism in the data. Quantitatively, we show that the interaction between selection and endogenous growth substantially amplifies the long-run output consequences of temporary shocks. The broader lesson is that the long-run effects of recessions depend not only on the fall in aggregate innovation, but also on the cross-sectional pattern of firms' innovation responses.

The rest of the paper is organized as follows. Section 2 presents the motivating evidence on heterogeneity in firm-level R&D and its cyclicality. Section 3 develops the quantitative model. Section 4 isolates the selection mechanism theoretically and provides direct empirical evidence for its key prediction. Section 5 brings the model to the data. Section 6 studies the aggregate quantitative implications, decomposes the interaction between selection and endogenous growth, and evaluates the resulting output costs of a severe downturn. Section 7 concludes.

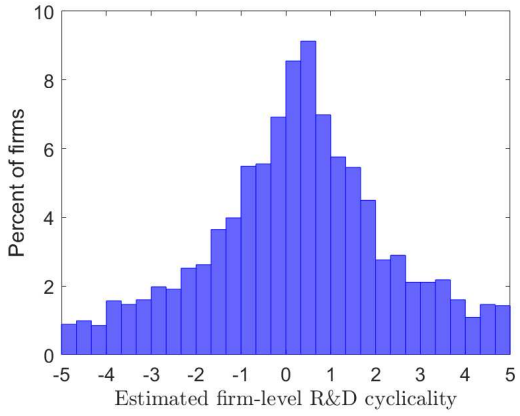
2 Motivating Evidence on Firm-Level R&D

We begin by documenting two features of firm-level R&D that discipline and motivate our quantitative analysis. The first concerns the cross-sectional allocation of R&D across firms: R&D intensity declines systematically with firm size. The second concerns dynamics: firms differ substantially in the responsiveness of their R&D expenditure to cyclical conditions. The first fact is well established in the literature, while the second is, to our knowledge, a new empirical fact and provides the central motivation for our analysis of heterogeneous R&D cyclicality.

Figure 1: Cross-sectional heterogeneity in R&D intensity and cyclicality.



(a) R&D intensity and firm size



(b) Firm-level R&D cyclicality

Notes: Subfigure 1a reports average log(1 + R&D/sales) and average log real sales within five equally populated real-sales groups. Subfigure 1b reports the distribution of estimated firm-level coefficients $\theta_{1,i}$ from regression Equation (1) of real R&D growth on industry output growth. Cyclicality coefficients are truncated at $[-5, 5]$.

The subfigures presented in Figure 1 use an unbalanced panel of Compustat firms over the period 1975–2018. Subfigure 1a documents the relationship between firm size and R&D intensity. We divide firm-year observations into five equally populated groups according to real sales, our measure of firm size, expressed in 1975 dollars as in Ignaszak and Sedláček (2026). The figure reports average real sales and average R&D intensity, measured as log(1 + R&D/sales), within each group. Consistent with a well-established pattern in the literature, R&D intensity declines with firm size (Akcigit and Kerr, 2018). This matters for growth since smaller firms may generate innovation and productivity gains that are large relative to their current economic size. This relationship disciplines the cross-sectional allocation of R&D in the model and is a central target of the steady-state calibration discussed in Section 5.

We next turn from levels to dynamics and study the distribution of firm-level R&D cyclicality in Subfigure 1b. To this end, we supplement the Compustat data with BEA measures of industry-level gross output. For each firm i, we estimate a separate regression:

$$\Delta rd_{it} = a_i + \theta_{1,i} \Delta out_{j(i)t} + u_{it}, \tag{1}$$

where Δrd_{it} is growth in real R&D expenditure and $\Delta out_{j(i)t}$ is real output growth in firm i's industry, denoted $j(i)$. Parameter a_i is a firm-specific constant, $\theta_{1,i}$ measures the cyclicality of firm i's R&D expenditure with respect to conditions in its industry and u_{it} is a residual. A positive value of $\theta_{1,i}$ therefore indicates that the firm's R&D expenditure is procyclical over its observed lifetime in the sample.

Subfigure 1b shows that the distribution of estimated firm-level coefficients $\theta_{1,i}$ is

centered on positive values. Innovation spending is therefore procyclical on average, consistent with the aggregate evidence in the literature (e.g. Barlevy (2007)). The cross-section, however, is considerably richer than this average suggests. The distribution is unimodal but broad — firms differ in both magnitude and sign of their responsiveness. The positive average response therefore does not describe a common firm-level relationship between R&D and the cycle. Rather, aggregate procyclicality coexists with substantial cross-sectional heterogeneity in firms' cyclical R&D behavior. This heterogeneity in responsiveness is distinct from the well-known heterogeneity in R&D intensity documented in Subfigure 1a: firms differ both in how intensively they undertake R&D and in how that expenditure moves with economic conditions.

Taken together, the two panels characterize distinct but complementary dimensions of firm-level R&D behavior. Subfigure 1a describes the cross-sectional allocation of R&D across firms, while Subfigure 1b describes the heterogeneity in its response to cyclical conditions. These two dimensions interact in aggregation: a firm's contribution to growth cyclicality depends both on the intensity of R&D it undertakes and on the responsiveness of that expenditure to the cycle. As such, these empirical patterns motivate the key firm-level dimensions used to discipline the quantitative model developed below.

3 Model

3.1 Environment

The model is in discrete time $t \in \{0, 1, 2, ...\}$ and features three different agents — intermediate goods firms, households and final goods producers. The latter two agents are representative and standard, while the interesting choices and dynamics are at the level of heterogeneous intermediate producers.

Household. The representative household inelastically supplies labor, N, to intermediate goods producers and chooses consumption, C, and asset holdings, A, to maximize life-time utility

$$\max_{\{C_t, A_{t+1}\}} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\sigma}}{1-\sigma} \tag{2}$$

where $\mathbb{E}$ is an expectations operator, $\beta \in [0, 1]$ is the discount factor and σ is the coefficient of relative risk aversion. The household makes its choices subject to the following budget constraint

$$P_t C_t + A_{t+1} = W_t N + (1 + R_t) A_t \tag{3}$$

where P_t is the aggregate price index, W_t is the wage rate and R_t is the return on assets.

Final goods firms. We assume that the final goods firm combines intermediate goods according to the following CES production function

$$Y_t = \left(\int_{j \in \Omega_t} y_t(j)^{\frac{\eta-1}{\eta}} dj \right)^{\frac{\eta}{\eta-1}} \tag{4}$$

where $y_t(j)$ denotes the quantity of intermediate good j at time t, Ω_t is both the measure and set of incumbent intermediate firms and $\eta > 1$ is the elasticity of substitution between varieties. These final goods producers are assumed perfectly competitive, taking the CPI P_t as given.

Incumbent intermediate goods firms. Each intermediate variety is associated with a firm-specific technology that can only be operated by a single firm. As a result, productivity and innovation outcomes are embodied within individual firms and are not transferable across producers. These firms operate with the objective of maximizing the expected present value of dividends to the firm's equityholders using production function

$$y_t(j) = Z_t q_t(j) n_t(j), \tag{5}$$

where $n_t(j)$ denotes labor resources utilized, $q_t(j)$ is *endogenous* firm-level productivity, which is driven by firms' investment into research and development (R&D) and Z_t is exogenous aggregate total factor productivity. The latter is assumed to evolve according to

$$\log(Z_t) = \rho_Z \log(Z_{t-1}) + \epsilon_{Z,t} \tag{6}$$

where $\epsilon_{Z,t} \sim N(0, \sigma_Z^2)$ is a stochastic shock. Intermediate goods firms engage in monopolistic competition, setting prices as a constant markup of $\eta/(\eta-1)$ over marginal cost. They hire from a perfectly competitive labor market. There are also two potential sources of exit — exogenous at rate $\delta \in [0,1]$ and endogenous due to the presence of fixed costs. These fixed costs ϕ are drawn i.i.d. each period, their level costs in terms of final goods are denoted $Q_t \phi$; we denote the distribution of normalized fixed costs as $F(\phi)$.

Firms purchase $Q_t s_t(j)$ units of the final good for R&D, where $s_t(j)$ denotes R&D expenditure relative to average productivity. To attain a probability of successful innovation $x_t(j) \in [0,1]$, a firm incurs normalized R&D expenditure

$$s_t(j) = \gamma \widehat{q}_t(j)^{\chi} x_t(j)^{\psi}, \tag{7}$$

where parameter $\psi > 1$ controls the curvature, $\gamma > 0$ is a shifter and $\chi > (\eta - 1)$

gives efficiency that varies with relative productivity $\widehat{q}_t(j) = q_t(j)/Q_t$, where average productivity is given by

$$Q_t = \left(\frac{1}{\Omega_t} \int_{j \in \Omega_t} q_t(j)^{\eta-1} dj \right)^{\frac{1}{\eta-1}}. \tag{8}$$

Equation (7) specifies that R&D costs are convex in the chosen probability of success and are an increasing function of a firm's relative productivity. In the event of successful innovation, a firm receives an improvement to its idiosyncratic productivity level, where the gross innovation step size $1+\lambda$ is drawn from an i.i.d. exogenous distribution $H(\lambda)$. In the absence of successful innovation, a firm's productivity remains unchanged. As average productivity in the economy grows, the firm's relative productivity therefore declines at the rate of average productivity growth, $G_t = Q_t/Q_{t-1} - 1$.

Entrant intermediate goods firms. Entrants incur a sunk cost in levels of final goods $Q_t \phi_e > 0$ at establishment in period t. Subsequently, they develop a prototype variety, whose idiosyncratic productivity $\widehat{q}_e(j)$ is drawn from a distribution $\omega_{e,t}(\widehat{q}_e) = \omega_t(\widehat{q})/\Omega_t$, where $\omega_t(\widehat{q})$ denotes the current period incumbent productivity measure and Ω_t is the incumbent mass. Similarly to the standard Schumpeterian growth model, these knowledge spillovers at entry sustain long-run growth. After drawing prototype productivity $\widehat{q}_{e,t}(j)$, an entrant chooses its probability of successful innovation $x_{e,t}(j)$. Its normalized R&D expenditure is

$$s_{e,t}(j) = \gamma_e \widehat{q}_e(j)^{\chi} x_{e,t}(j)^{\psi}. \tag{9}$$

Thus, the entrant purchases $Q_t s_{e,t}(j)$ units of the final good for R&D. Parameter γ_e is an entrant-specific shifter and parameters $\psi > 1$ and $\chi > \eta - 1$ are as in the incumbent technology Equation (7). We assume that these potential entrants must successfully innovate upon their prototype in order to become incumbents at time $t+1$. When successful, they draw a step size from the distribution $H(\lambda)$ and commence activities as an incumbent firm with relative productivity $(1+\lambda)\widehat{q}_e/(1+G)$ at $t+1$. Should they fail to innovate, they exit the industry and their prototype becomes obsolete. We denote the mass of potential entrants at t by $\Omega_{e,t}$.

Normalization. Variables that inherit the economy's productivity trend are expressed in units of Q_t when we formulate the stationary firm problem, in particular

$$\widehat{q}_t(j) \equiv \frac{q_t(j)}{Q_t}, \qquad \widehat{y}_t(j) \equiv \frac{y_t(j)}{Q_t}, \qquad \widehat{Y}_t \equiv \frac{Y_t}{Q_t}, \qquad \widehat{W}_t \equiv \frac{W_t}{Q_t}, \qquad \widehat{C}_t \equiv \frac{C_t}{Q_t}.$$

To avoid overloading the notation with hats, lower-case $s_t(j)$ and $s_{e,t}(j)$ denote incumbent and entrant R&D expenditure in units of Q_t. Likewise, fixed operating costs ϕ and the sunk entry cost ϕ_e are expressed in units of Q_t. When introduced below, profits π and firm values v, $\widetilde{v}$, and v_e are also expressed in units of Q_t. The distribution F is defined over normalized fixed costs.

Timing. The period t starts with the activities of incumbents.

1. The aggregate TFP level Z_t is realized from Equation (6).
2. Draw i.i.d. fixed cost shock ϕ from $F(\phi)$ and make endogenous exit choice.
3. Conditional on remaining, make optimal static choices of $n_t(j), y_t(j)$ in accordance with technology (5).
4. Make optimal inter-temporal R&D choice $s_t(j), x_t(j)$ through technology (7) and distribute dividends to households.
5. Draw exogenous death shock.
6. Draw innovation success shock from Bernoulli$(x_t(j))$, as well as step size from $H(\lambda)$ and move to time $t+1$.

Then the new entrants move with timing at t as follows.

i. Pay sunk cost ϕ_e to create variety j and draw prototype productivity level.

ii. Make optimal choice of R&D investment $s_{e,t}(j)$ and $x_{e,t}(j)$.

iii. Draw innovation success shock from Bernoulli$(x_{e,t}(j))$ and step size from $H(\lambda)$.

iv. Become an incumbent at $t+1$ if successful in innovating, otherwise exit.

3.2 Equilibrium

Here we detail the optimal choices of agents, given the setup in Section 3.1.

Household. The household prices claims to the intermediate goods firms' cash flows using Euler equation

$$1 = \beta \mathbb{E}_t \left[(1 + R_{t+1}) \left(\frac{C_{t+1}}{C_t} \right)^{-\sigma} \right], \tag{10}$$

for an arbitrary return $R_{t+1} \geqslant 0$. Equation (10) determines the stochastic discount factor used to price claims to firms' cash flows.

Final goods firm. The final goods firm chooses its purchases of variety j to maximize its profits, subject to aggregation technology (4), giving variety-level demand

$$y_t(j) = p_j(t)^{-\eta} Y_t, \tag{11}$$

as well as the definition of the CPI

$$P_t = \left(\int_{j \in \Omega_t} p_t(j)^{1-\eta} dj \right)^{\frac{1}{1-\eta}}. \tag{12}$$

We take the final good to be the numéraire, meaning that $P_t = 1 \, \forall t$.

Incumbent intermediate goods firms. In what follows, we omit variety-level notation and instead switch to describing firms by their state. We also omit t subscripts, where future values are denoted by $'$ superscripts. Firms' choices are contingent on their idiosyncratic state $\widehat{q}$ as well as the aggregate state, which we denote as Γ, comprising the TFP shock Z and firm cross-sectional distribution $\omega(\widehat{q})$. We then write the Bellman equation, conditional on remaining in operation, as

$$v(\widehat{q}, \Gamma) = \max_{p,n,x} \pi(\widehat{q}, \Gamma) + (1-\delta) \mathbb{E} \widehat{M}' \left\{ x \int \widetilde{v} \left(\frac{\widehat{q}(1+\lambda)}{1+G'}, \Gamma' \right) dH(\lambda) + (1-x) \widetilde{v} \left(\frac{\widehat{q}}{1+G'}, \Gamma' \right) \right\}. \tag{13}$$

The first term in (13), $\pi(\widehat{q}, \Gamma)$, denotes the firm's current-period profits. The second term captures the expected continuation value, conditional on surviving the exogenous death shock, which occurs with probability $1 - \delta$. With probability x, the firm successfully innovates and draws a productivity improvement $1 + \lambda$ from distribution $H(\lambda)$. In this case, next period's relative productivity is given by $\widehat{q}(1+\lambda)/(1+G')$, reflecting both the innovation gain and the erosion of relative productivity due to aggregate productivity growth. With probability $1 - x$, innovation is unsuccessful and the firm's relative productivity evolves to $\widehat{q}/(1+G')$. The value function $\widetilde{v}$ denotes the firm's equity value at the beginning of the following period, prior to the realization of the fixed-cost shock and the associated endogenous exit decision. The firm maximization is subject to the stationary SDF $\widehat{M}'$ associated with (10), the demand curve (11), innovation technology (7), production function (5) and TFP process (6).

Period profits and the start-of-period equity value are given by

$$\pi(\widehat{q}, \Gamma) = p\widehat{y} - \widehat{W} n - s \tag{14}$$

$$\widetilde{v}(\widehat{q}, \Gamma) = \int_{\phi} \max \left[v(\widehat{q}, \Gamma) - \phi, 0 \right] dF(\phi), \tag{15}$$

Equation (14) states that firm profits consist of revenues net of labor costs and R&D expenditures. Equation (15) gives the firm's equity value at the beginning of the period, after the aggregate state has been realized but prior to the realization of its stochastic fixed-cost shock and the associated endogenous exit decision. Upon drawing ϕ, the firm continues operating if $v(\widehat{q},\Gamma)-\phi \geqslant 0$, and exits otherwise. Limited liability implies that continuation values are bounded below by zero. This implies the existence of a cutoff fixed-cost draw, denoted $\widetilde{\phi}(\widehat{q},\Gamma)=v(\widehat{q},\Gamma)$, such that a firm exits whenever $\phi>\widetilde{\phi}(\widehat{q},\Gamma)$ and continues operating otherwise. The resulting optimal decisions are given by

$$\begin{aligned} p(\widehat{q},\Gamma) &= \frac{\eta}{\eta-1}\frac{\widehat{W}}{Z\widehat{q}}, \\ n(\widehat{q},\Gamma) &= \frac{\widehat{y}}{Z\widehat{q}} \\ \widehat{y}(\widehat{q},\Gamma) &= p(\widehat{q},\Gamma)^{-\eta}\widehat{Y} \\ x(\widehat{q},\Gamma) &= \left[\frac{(1-\delta)\mathbb{E}\widehat{M}'\Delta(\widehat{q},\Gamma')}{\gamma\psi\widehat{q}^{\chi}}\right]^{\frac{1}{\psi-1}} \end{aligned} \tag{16}$$

where $p(\widehat{q},\Gamma)$, $n(\widehat{q},\Gamma)$, $\widehat{y}(\widehat{q},\Gamma)$ and $x(\widehat{q},\Gamma)$ denote the firm's optimal price, labor demand, output and innovation effort policies, respectively. The policy function for the innovation choice takes the transformed discounted value of the marginal benefit to R&D effort relative to the marginal cost, where the latter comprises terms from the R&D technology (7) and the former is given by

$$\Delta(\widehat{q},\Gamma')\equiv\int\left[\widetilde{v}\left(\frac{\widehat{q}(1+\lambda)}{1+G'},\Gamma'\right)-\widetilde{v}\left(\frac{\widehat{q}}{1+G'},\Gamma'\right)\right]dH(\lambda),$$

which represents the expected continuation-value gain from R&D success, integrating over the distribution $H(\lambda)$ of stochastic step sizes.

Entrant intermediate goods firms. Entrants make a choice of R&D investment, conditional on their prototype draw, to maximize conditional value

$$v_e(\widehat{q}_e,\Gamma)=\max_{\{x_e\}}-s_e+\mathbb{E}\widehat{M}'x_e\Delta_e(\widehat{q}_e,\Gamma'), \tag{17}$$

subject to their R&D technology (9); the benefit of innovating is given by

$$\Delta_e(\widehat{q}_e,\Gamma')\equiv\int\widetilde{v}\left(\frac{\widehat{q}_e(1+\lambda)}{1+G'},\Gamma'\right)dH(\lambda).$$

The problem in (17) then yields optimal innovation choice

$$x_e(\widehat{q}_e, \Gamma) = \left[\frac{\mathbb{E}\widehat{M}' \Delta_e(\widehat{q}_e, \Gamma')}{\gamma_e \psi \widehat{q}_e^{\chi}} \right]^{\frac{1}{\psi - 1}}. \tag{18}$$

Since incumbency is contingent on successful innovation, the benefit of R&D is given by the absolute continuation value $\widetilde{v}$, rather than by the incremental value gain $\Delta(\widehat{q}, \Gamma')$ relevant for incumbents. Using the recursive formulation in (17) and optimal R&D choice in (18), one can define the ex-ante value to entry as

$$\widetilde{v}_e(\Gamma) = -\phi_e + \mathbb{E}_e[v_e(\widehat{q}_e, \Gamma)], \tag{19}$$

which represents the expected entry value net of the sunk cost of creating a new variety. Note that the expectation operator is taken with respect to the draws of prototype relative productivity from the incumbent cross-sectional distribution.

Equilibrium definition. We define equilibrium such that the following list of conditions hold.

1. All agents are optimizing.
2. All markets are clearing. For the final-goods market, aggregate output is allocated to household consumption, R&D expenditures by incumbent firms and potential entrants, incumbent operating fixed costs, and sunk entry costs:

$$\widehat{Y} = \widehat{C} + S_I + S_E + \Phi_I + \Phi_E, \tag{20}$$

where

$$\begin{aligned}
S_I &= \int F\left(\widetilde{\phi}(\widehat{q}, \Gamma)\right) s(\widehat{q}, \Gamma)\, \omega(\widehat{q}, Z)\, d\widehat{q}, \\
S_E &= \Omega_e \int s_e(\widehat{q}_e, \Gamma) \omega_e(\widehat{q}_e, Z)\, d\widehat{q}_e, \\
\Phi_I &= \int \left[\int_0^{\widetilde{\phi}(\widehat{q}, \Gamma)} \phi\, dF(\phi) \right] \omega(\widehat{q}, Z)\, d\widehat{q}, \\
\Phi_E &= \Omega_e \phi_e.
\end{aligned}$$

Here, $\omega(\widehat{q}, Z)$ denotes the cross-sectional measure of incumbent firms, so that integration over ω already aggregates across the incumbent firm mass. Incumbent R&D expenditure is incurred only by firms whose fixed-cost draw lies below the endogenous continuation cutoff $\widetilde{\phi}(\widehat{q}, \Gamma)$, while Φ_I collects the corresponding operating fixed costs. Potential entrants pay the sunk entry cost ϕ_e and undertake R&D

before learning whether innovation is successful, implying aggregate entrant expenditures S_E and Φ_E as defined above. All terms in (20) are measured in units of Q_t. Thus, S_I and S_E are aggregate normalized R&D expenditures, while Φ_I and Φ_E are aggregate normalized operating and entry costs. Multiplying Equation (20) by Q_t recovers the resource constraint in levels.

The labor market clears according to

$$L = N. \tag{21}$$

which says that aggregate labor demand from firms L equals the inelastic supply coming from households N.

3. The law of motion for the firm cross section $\omega(\widehat{q}, Z)$ is satisfied

$$\begin{aligned}\omega'(\widehat{q}, Z') =&(1-\delta)\int^{\widetilde{\phi}(\widehat{q},\Gamma')}\int_{\lambda} x\left(\frac{1+G'}{1+\lambda}\widehat{q},\Gamma\right)\omega\left(\frac{1+G'}{1+\lambda}\widehat{q},Z\right)\mathrm{d}F(\phi)\mathrm{d}H(\lambda)\\ &+(1-\delta)\int^{\widetilde{\phi}(\widehat{q},\Gamma')}\left\{1-x\left([1+G']\widehat{q},\Gamma\right)\right\}\omega\left([1+G']\widehat{q},Z\right)\mathrm{d}F(\phi)\\ &+\Omega_e\int^{\widetilde{\phi}(\widehat{q},\Gamma')}\int_{\lambda} x_e\left(\frac{1+G'}{1+\lambda}\widehat{q},\Gamma\right)\omega_e\left(\frac{1+G'}{1+\lambda}\widehat{q},Z\right)\mathrm{d}F(\phi)\mathrm{d}H(\lambda)\end{aligned} \tag{22}$$

where the first term denotes the inflow of surviving incumbent firms that reach state $\widehat{q}$ through successful innovation. The second term denotes the inflow of surviving incumbent firms that fail to innovate and therefore drift down the relative productivity distribution due to aggregate productivity growth. The final term captures the inflow of successful entrants that become incumbent firms.

4. The entry mass Ω_e is determined from the free entry condition $\widetilde{v}_e(\Gamma) = 0$, using the expression in (19).

5. We define the growth rate of the average productivity level in the economy, drawing on the cross-sectional measure law of motion (22). We denote the growth rate as $G' = Q'/Q - 1$. From the definition of Q in (8), we can define the growth rate in variety-level notation as

$$(1+G')^{\eta-1} = \frac{1}{\Omega'}\int_{j\in\Omega'}\left(\frac{q'}{Q}\right)^{\eta-1} dj. \tag{23}$$

Leveraging the cross-section (22) and re-expressing (23) in terms of state variables

$$\begin{aligned}\Omega'(1+G')^{\eta-1} &= (1-\delta)\int_{\widehat{q}}\int^{\widetilde{\phi}(\widehat{q}/(1+G'),\Gamma')} \widehat{q}^{\eta-1}\left[1-x(\widehat{q},\Gamma)\right]\omega(\widehat{q},Z)\; dF(\phi)\, d\widehat{q} \\ &+ (1-\delta)\int_{\widehat{q}}\int_{\lambda}\int^{\widetilde{\phi}(\widehat{q}(1+\lambda)/(1+G'),\Gamma')} \left[\widehat{q}(1+\lambda)\right]^{\eta-1} x(\widehat{q},\Gamma)\,\omega(\widehat{q},Z)\; dF(\phi)\, dH(\lambda)\, d\widehat{q} \\ &+ \Omega_e\int_{\widehat{q}_e}\int_{\lambda}\int^{\widetilde{\phi}(\widehat{q}_e(1+\lambda)/(1+G'),\Gamma')} \left[\widehat{q}_e(1+\lambda)\right]^{\eta-1} x_e(\widehat{q}_e,\Gamma)\,\omega_e(\widehat{q}_e,Z)\; dF(\phi)\, dH(\lambda)\, d\widehat{q}_e,\end{aligned} \tag{24}$$

which correspond to the contributions of unsuccessful incumbent firms, successful incumbent firms and successful entrants, respectively.

4 Mechanism

Having detailed the full quantitative model, we now isolate the main economic mechanism through a simplified version of the framework. This exercise yields a set of empirical predictions that we subsequently evaluate using firm-level micro data.

4.1 Selection and Heterogeneous R&D Cyclicality

We now isolate the mechanism generating heterogeneous firm-level R&D responses to aggregate fluctuations. The key result is that such heterogeneity does not require firm-specific shocks, financial frictions, or asymmetric adjustment costs. Under CES demand, aggregate shocks generate the same first-order percentage change in gross cashflows for all firms. In the absence of endogenous exit, firms therefore have identical R&D elasticities to aggregate conditions. Heterogeneous R&D cyclicality arises endogenously because selection changes the sensitivity of the value gain from successful innovation to aggregate cashflows.

To show this result transparently, consider a two-period version of the model, with $t \in \{0,1\}$. We assume there is no exogenous exit ($\delta = 0$), and the innovation step-size distribution is degenerate, with $H(\bar{\lambda}) = 1$ for a deterministic $\bar{\lambda}$. We consider the problem of an incumbent at time $t = 0$, following the realization of the aggregate shock Z and conditional on having survived its current fixed-cost draw. Using the optimal static choices in equation (16), period cashflow gross of the fixed operating cost can be written as

$$\pi = \left(\frac{\widehat{W}}{Z}\right)^{1-\eta}\left(\frac{\eta}{\eta-1}\right)^{-\eta}\frac{1}{\eta-1}\widehat{Y}\,\widehat{q}^{\eta-1} - s. \tag{25}$$

We suppress the dependence of policy and value functions on the firm's state variables

to simplify notation.

The role of selection can already be seen from terminal-period equity value. Since firms undertake no further R&D at $t = 1$, their value conditional on operating is simply gross operating cashflow,

$$v' = \pi'.$$

Prior to observing the terminal fixed-cost draw, however, equity value is

$$\widetilde{v}'(\pi') = \int_0^{\pi'} (\pi' - \phi) dF(\phi).$$

Its marginal value is

$$\frac{\partial \widetilde{v}'(\pi')}{\partial \pi'} = F(\pi'), \qquad \frac{\partial^2 \widetilde{v}'(\pi')}{\partial (\pi')^2} = f(\pi') > 0.$$

The first expression has a simple interpretation. An additional unit of cashflow is valuable only when the firm remains active, so its marginal value equals the probability of survival. Since survival becomes more likely as cashflow rises, the marginal value of cashflow rises as well. Endogenous exit therefore makes $\widetilde{v}'$ convex in terminal operating cashflow. Without endogenous exit, survival is certain and $\widetilde{v}'$ is linear in cashflow.

Let $+$ and $-$ denote the terminal-period states following successful and unsuccessful innovation, respectively. The time-zero problem can then be written as

$$\begin{aligned} v &= \max_s \left\{ \pi + \mathbb{E}\widehat{M}' \left[x\widetilde{v}'_+ + (1-x)\widetilde{v}'_- \right] \right\} \\ &= \max_s \left\{ \pi + \mathbb{E}\widehat{M}' \left[xF'_+ \left(\pi'_+ - c'_+ \right) + (1-x)F'_- \left(\pi'_- - c'_- \right) \right] \right\}, \end{aligned} \tag{26}$$

where

$$F'_\pm \equiv F(\pi'_\pm)$$

are the corresponding survival probabilities and

$$c'_\pm \equiv \frac{1}{F'_\pm} \int_0^{\pi'_\pm} \phi \, dF(\phi)$$

are expected fixed costs conditional on continuation. Successful innovation raises terminal productivity and therefore $\pi'_+ > \pi'_-$. Whenever the firm is exposed to selection, successful innovation also raises its probability of survival.

Using the R&D technology $s = \gamma \widehat{q}^{\chi} x^{\psi}$, the first-order condition for innovation gives

$$x = \left[\frac{\mathbb{E}\widehat{M}' \left[F'_+ \left(\pi'_+ - c'_+ \right) - F'_- \left(\pi'_- - c'_- \right) \right]}{\psi \gamma \widehat{q}^{\chi}} \right]^{\frac{1}{\psi - 1}}. \tag{27}$$

Innovation is therefore determined by the difference between terminal equity values following successful and unsuccessful innovation, $\widetilde{v}'_+ - \widetilde{v}'_-$.

We next take a first-order approximation of Equation (27) around the deterministic steady state. Appendix A provides the derivation. Let bars denote steady-state values and checks percentage deviations from steady state. The innovation response can be written as

$$\begin{aligned}
\check{x} &= \frac{1}{\psi - 1}\mathbb{E}\left\{\check{M}' + \kappa\check{\pi}'\right\}, \\
\check{\pi}' &= -(\eta-1)\check{W}' + (\eta-1)\check{Z}' + \check{Y}' - (\eta-1)\frac{\bar{G}}{1+\bar{G}}\check{G}', \\
\kappa &= \frac{\bar{F}'_+\bar{\pi}'_+ - \bar{F}'_-\bar{\pi}'_-}{\bar{F}'_+(\bar{\pi}'_+ - \bar{c}'_+) - \bar{F}'_-(\bar{\pi}'_- - \bar{c}'_-)}.
\end{aligned} \tag{28}$$

Note that expression $\check{M}'$ captures movements in the SDF and is common across firms. The CES structure has implications for the cashflow term $\kappa\check{\pi}'$. Successful and unsuccessful innovators differ in the *level* of their future cashflows, but their percentage response to aggregate variables is identical. Under CES demand and constant markups, relative productivity enters gross cashflows multiplicatively, so that its level drops out of $\check{\pi}'$. Hence, absent selection, heterogeneous firms would have the same first-order R&D elasticity to aggregate conditions. All heterogeneity in their first-order innovation responses therefore operates through the loading κ.

The economic meaning of κ can be seen directly from the value gain from successful innovation. Let a denote a common proportional shifter of terminal cashflows, so that

$$\pi'_\pm(a) = a\bar{\pi}'_\pm,$$

and define

$$\Delta(a) \equiv \widetilde{v}'_+(a) - \widetilde{v}'_-(a),$$

where

$$\widetilde{v}'_\pm(a) = \int_0^{a\bar{\pi}'_\pm} \left(a\bar{\pi}'_\pm - \phi\right) dF(\phi).$$

At the deterministic steady state, $a = 1$, and

$$\Delta(1) = \bar{F}'_+(\bar{\pi}'_+ - \bar{c}'_+) - \bar{F}'_-(\bar{\pi}'_- - \bar{c}'_-),$$

which is exactly the denominator of κ in Equation (28). Moreover,

$$\frac{d\widetilde{v}'_\pm(a)}{da} = F(a\bar{\pi}'_\pm)\bar{\pi}'_\pm,$$

and hence

$$\left.\frac{d\Delta(a)}{da}\right|_{a=1} = \bar{F}'_+\bar{\pi}'_+ - \bar{F}'_-\bar{\pi}'_-,$$

which is exactly its numerator. It follows that

$$\kappa = \frac{\left.\frac{d\Delta(a)}{da}\right|_{a=1}}{\Delta(1)} = \left.\frac{d\log\Delta(a)}{d\log a}\right|_{a=1}.$$

Thus, κ is the elasticity of the terminal value gain from successful innovation, $\tilde{v}'_+ - \tilde{v}'_-$, with respect to the common proportional movement in future cashflows induced by the aggregate shock.

Proposition 1. *Absent endogenous exit, CES demand implies that all firms have the same innovation elasticity to an aggregate shock. When innovation affects survival, endogenous selection raises this elasticity, so that* $\kappa > 1$.

To see the result, note that

$$\left.\frac{d\Delta(a)}{da}\right|_{a=1} - \Delta(1) = \bar{F}'_+\bar{c}'_+ - \bar{F}'_-\bar{c}'_- = \int_{\bar{\pi}'_-}^{\bar{\pi}'_+} \phi\, dF(\phi).$$

The interval $(\bar{\pi}'_-, \bar{\pi}'_+]$ has a direct economic interpretation: it contains the fixed-cost realizations for which innovation changes the firm's survival decision. If innovation fails, the firm exits because $\phi > \bar{\pi}'_-$; if it succeeds, the higher cashflow $\bar{\pi}'_+$ allows the firm to continue. Whenever the fixed-cost distribution assigns positive probability to this interval, the integral is strictly positive. Hence

$$\left.\frac{d\Delta(a)}{da}\right|_{a=1} > \Delta(1),$$

and therefore $\kappa > 1$. Endogenous selection thus makes the value gain from successful innovation respond more than proportionally to a common movement in future cashflows.

Without endogenous exit, survival is certain following either innovation outcome, so $\bar{F}'_+ = \bar{F}'_- = 1$ and $\bar{c}'_+ = \bar{c}'_-$. More generally, if successful innovation does not change the probability of survival, then there are no fixed-cost realizations for which success makes the difference between exit and continuation. The set of surviving fixed-cost draws is therefore the same after success and failure, and so is their conditional mean.[3] It follows that

$$\left.\frac{d\Delta(a)}{da}\right|_{a=1} = \Delta(1),$$

[3] If $\bar{F}'_+ = \bar{F}'_-$, then

$$\bar{F}'_+ - \bar{F}'_- = \Pr(\bar{\pi}'_- < \phi \leqslant \bar{\pi}'_+) = 0.$$

Hence the fixed-cost distribution places no probability mass between the two survival cutoffs. Splitting

and hence $\kappa = 1$.

The intuition is straightforward. Without endogenous exit, successful innovation raises future profitability but does not affect survival. A common proportional movement in future cashflows therefore changes the value gain from innovation proportionally. With selection, successful innovation also moves the firm to a state with a higher probability of survival. The value gain from innovation then responds more than proportionally to aggregate conditions: following an expansion, firms exposed to selection raise R&D more strongly, while following a contraction they cut it more strongly.

The convexity of $\widetilde{v}'$ provides the mathematical counterpart to this mechanism. Since $\frac{\partial \widetilde{v}'}{\partial \pi'} = F(\pi')$, the marginal value of cashflow rises with the probability of survival. Successful innovation therefore places the firm at a point where an additional increase in cashflow is more valuable than after unsuccessful innovation. This is the source of the selection amplification, $\kappa > 1$.

Proposition 1 concerns heterogeneity in innovation *elasticities*, not heterogeneity in innovation levels. Firms may choose different steady-state levels of R&D because innovation costs vary with relative productivity through $\widehat{q}^{\chi}$.[4] Endogenous selection can further contribute to heterogeneity in innovation levels by making the value gain from successful innovation depend on a firm's position relative to the exit margin. The key result of Proposition 1, however, is that heterogeneity in innovation levels alone does not generate heterogeneous cyclicality: absent selection, firms have the same percentage response of innovation to an aggregate shock even when their steady-state innovation choices differ. Endogenous selection is what breaks this neutrality and generates heterogeneous R&D elasticities.

We next push the insight a step further and ask how the selection amplification identified in Proposition 1 varies across the productivity distribution. Proposition 1 establishes that selection amplifies innovation responses, but by itself does not imply how this amplification varies across firms. We can establish this relationship analytically under a sufficient restriction on the fixed-cost distribution.

Proposition 2. *Suppose that the CDF of log fixed costs, $\ell = \log(\phi) \mapsto F(e^{\ell})$, is log-concave. Then the loading κ is non-increasing in relative productivity $\widehat{q}$. Hence lower-productivity firms have larger first-order innovation elasticities to aggregate conditions.*

Proposition 2 provides a sufficient condition under which the selection amplification

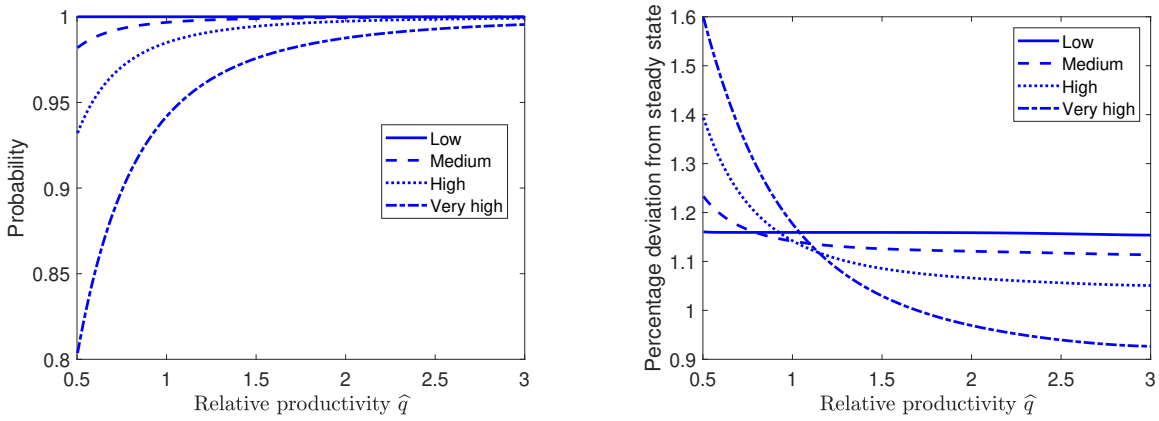
the larger integral,

$$\int_0^{\bar{\pi}'_+} \phi \, dF(\phi) = \int_0^{\bar{\pi}'_-} \phi \, dF(\phi) + \int_{\bar{\pi}'_-}^{\bar{\pi}'_+} \phi \, dF(\phi) = \int_0^{\bar{\pi}'_-} \phi \, dF(\phi).$$

Together with $\bar{F}'_+ = \bar{F}'_-$, this implies $\bar{c}'_+ = \bar{c}'_-$.

[4] Absent endogenous selection, equation (27) implies heterogeneous steady-state innovation levels whenever $\chi \neq \eta - 1$. When $\chi = \eta - 1$, firms choose the same steady-state level of x.

Figure 2: Numerical example with varying average fixed costs.



(a) Survival probabilities

(b) Impact responses

Notes: The panels report results for low, medium, high and very high average fixed costs. Subfigure 2a depicts steady-state survival probabilities. Subfigure 2b shows the impact response, corresponding to the first period of the IRF, of optimal R&D choices $x(\widehat{q}, \Gamma)$. Percentage deviations are multiplied by 100.

declines monotonically with productivity. The intuition is closely related to Proposition 1. Lower-productivity firms have lower operating cashflows and are closer to the exit margin. For these firms, successful innovation can materially change the probability of continuation: failure leaves the firm relatively exposed to exit, while success moves it to a higher-cashflow state with a higher probability of survival. Their innovation incentives are therefore particularly sensitive to aggregate conditions.

For highly productive firms, by contrast, survival is close to certain following either innovation outcome. Successful innovation still raises profits, but has relatively little effect on survival. As productivity rises, this additional survival channel becomes progressively weaker and κ approaches its no-selection benchmark of one. The log-concavity assumption provides a convenient sufficient condition ensuring that this decline is monotonic.[5]

This result gives the cross-sectional implication of the mechanism for R&D cyclicality. Following a positive aggregate shock, firms closer to exit raise R&D more strongly because successful innovation moves them to a substantially safer survival position, making the value gain from innovation particularly responsive to the improvement in aggregate cashflows. For firms already far from exit, innovation has little effect on survival, so this additional amplification is weak. Following a negative aggregate shock, the mechanism works in reverse: firms closer to exit experience a larger decline in the value gain from successful innovation and therefore cut R&D more strongly.

Similar logic carries over to the full infinite-horizon quantitative model described in Section 3, looking at the expression for $x(\widehat{q}, \Gamma)$ in equation (16). The difference in the full

[5]The condition is satisfied by a broad class of commonly used distributions. In particular, a log-concave density for log fixed costs implies a log-concave CDF. Examples include the normal, logistic, Laplace, Gumbel and uniform distributions. The lognormal fixed-cost distribution used in our quantitative model therefore satisfies the condition.

model is that the value function contains information regarding how innovation shapes survival probabilities in all future periods, as well as affecting the cost of R&D through the technology in Equation (7). Persistence in the TFP process ρ_Z therefore also plays a role, as it governs how strongly a current shock affects the continuation-value gain from improving $\widehat{q}$.

Figure 2 gives a numerical illustration of the mechanism in the full model. It shows steady-state survival probabilities (subfigure 2a) and the cross-section of impact responses of the R&D choice (subfigure 2b) following a 2% aggregate TFP shock. We depict these objects for four configurations of the fixed-cost distribution, increasing from low mean through very high. Under the lowest fixed cost, survival probabilities are numerically close to one for all firms, meaning that firms choose nearly identical pro-cyclical R&D responses to the shock.[6] As the fixed-cost burden rises, low-relative-productivity firms become increasingly exposed to endogenous exit, while high-relative-productivity firms remain largely insulated. The impact responses exhibit the opposite gradient: firms with the lowest survival probabilities are the most pro-cyclical. This is precisely the cross-sectional implication of Proposition 2: the amplification generated by selection is strongest among firms closer to the exit margin.

Taking stock. The results in this section establish the theoretical mechanism produced by the model and highlight its novelty. We establish a neutrality result for innovation in heterogeneous-firm models: in an otherwise standard environment with CES demand, heterogeneity in productivity, size, profitability, and steady-state R&D is not sufficient to generate heterogeneous first-order innovation responses to a common aggregate shock. Absent endogenous selection, all firms have the same innovation elasticity. Selection breaks this neutrality because successful innovation affects survival for firms close to the exit margin, making the return to innovation more sensitive to aggregate conditions.

This result distinguishes our mechanism from other heterogeneous-firm models of innovation over the business cycle. In Ates and Saffie (2021), financial crises increase the cost of financing new ideas to enter the market, which reduces the mass of new entrants but increase their quality. This quality-quantity trade off can magnify or dampen the effect of financial crisis. In Schmitz (2021), there is no selection and heterogeneous innovation responses to financial crisis arise because small firms are more exposed to financing constraints. These mechanisms add asymmetric financial exposure to the standard heterogeneous-firm innovation framework (Klette and Kortum, 2004). Our result instead shows that no firm-specific shock or asymmetric financial friction is required: in a heterogeneous-firm environment akin to Hopenhayn (1992), with innovation by incumbents and entrants in the spirit of Klette and Kortum (2004), heterogeneous R&D

[6]Note that the fixed-cost distribution $F(\phi)$ is parameterized as lognormal (see Section 5). As such, the survival probability is never exactly one.

cyclicality arises endogenously from the interaction between innovation and the exit margin. Proposition 1 shows why this margin is essential: without endogenous selection, the model delivers a common first-order R&D elasticity despite substantial underlying firm heterogeneity. The analytical predictions derived in the simplified two-period economy are confirmed by the numerical results from the full quantitative model, where firms closer to the exit margin display systematically stronger R&D responses to aggregate shocks.

4.2 Empirical Tests of the Selection Mechanism

Having highlighted that firm selection is the main driver of heterogeneity in R&D elasticities to the cycle in our model, we now confront this channel with data. In particular — we seek to test the prediction that firms closer to the exit margin are more procyclical with innovation. We return to the data of Section 2 — Compustat over 1975–2018 for information on firm-level R&D and BEA aggregate data. Our empirical design builds directly on the expressions in (28). In particular, note that the response of R&D spending can be written as $\check{s} = \psi/(\psi-1)\mathbb{E}\left\{\check{M}' + \kappa\check{\pi}'\right\}$, where recall steady state survival probabilities shape the loading through κ. We then take this to the data by running unbalanced regressions over firms i and time t of the form

$$\Delta rd_{it} = a_i + \theta_1 \Delta out_{j(i)t} + \theta_2 (sel_{it} \times \Delta out_{j(i)t}) + \theta_3 sel_{it} + u_{it} \tag{29}$$

where Δrd_{it} is growth in firm-level real R&D expenditures, a_i is a firm-level fixed effect, $\Delta out_{j(i)t}$ represents firm i's industry ($j(i)$) real output growth, sel_{it} is a (possibly time-varying) indicator variable capturing closeness to exit and u_{it} is a residual.[7] In this specification, industry output growth $\Delta out_{j(i)t}$ is our empirical proxy for the common cash-flow component in (28), while $sel_{it} \times \Delta out_{j(i)t}$ allows the firm-level R&D loading to vary with closeness to exit. As such, the parameter of interest is θ_2, which captures the incremental cyclicality of R&D associated with being near exit. In running these regressions, we consider varying combinations of time fixed effects and firm-level balance sheet control variables.[8] This exercise differs from the firm-level regressions in specification (1) in Section 2, which estimate a separate cyclicality coefficient for each firm in order to characterize the cross-sectional distribution of R&D responses. Here, we instead pool firm-year observations and ask whether R&D cyclicality varies systematically with firms' proximity to the exit margin, as predicted by the model.

We consider two measures of closeness to exit. The first is an ex-post measure: relating

[7] We use an unbalanced panel because restricting the sample to firms observed throughout would condition on continued presence in Compustat and mechanically exclude much of the firm turnover relevant for studying selection.

[8] Firm-level controls include current and lagged total assets and cash flow, total liabilities, gross property, plant and equipment, short-term debt and long-term debt, all normalized by sales. Cash flow is defined as income before extraordinary items plus depreciation.

Table 1: Selection mechanism regressions based on Specification (29).

	Exit-based measure				Distance-to-default measure			
Output growth θ_1	0.378***	0.557***	0.624**	0.782**	0.331***	0.477***	0.481**	0.576**
	(0.079)	(0.122)	(0.262)	(0.347)	(0.081)	(0.136)	(0.165)	(0.235)
Interaction θ_2	1.996***	2.543**	1.968***	2.479**	0.269**	0.312**	0.333***	0.402**
	(0.349)	(0.970)	(0.375)	(0.987)	(0.101)	(0.129)	(0.105)	(0.138)
Close to exit θ_3	−0.313***	−0.265***	−0.350***	−0.312***				
	(0.008)	(0.027)	(0.010)	(0.023)				
Time FE	Y	Y	N	N	Y	Y	N	N
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	N	Y	N	Y	N	Y	N
Observations	64,590	76,316	64,590	76,316	55,226	61,818	55,226	61,818

Notes: The dependent variable in all specifications is real R&D growth. Columns 1–4 use the exit-based measure and columns 5–8 use the distance-to-default measure. Standard errors are in parentheses. Standard errors clustered at the industry-level. Stars *, ** and *** denote significance at the 10%, 5% and 1% levels.

to firms that actually exit the sample. We restrict attention to those exiting for flagged reasons closest to exit in our model — bankruptcy or liquidation (Compustat DLRSN of 02 or 03). We then set $sel_{it} = 1$ when one such firm is in its final year before exit and to zero otherwise.

The second is an ex-ante measure: we construct firms' distance to default and implied probability of default. Following Corbae and D'Erasmo (2021), the one-year distance to default is

$$DD_{it} = \frac{\log(V_{it}/L_{it}) + \left(\mu_{A,i} - \frac{1}{2}\sigma^2_{A,i}\right)}{\sigma_{A,i}},$$

where V_{it} is the estimated market value of the firm's assets, L_{it} is short-term debt plus one-half of long-term debt, and $\mu_{A,i}$ and $\sigma_{A,i}$ are the mean and standard deviation of asset growth.[9] The implied probability of default is $PD_{it} = \Phi(-DD_{it})$, where Φ is the standard normal cumulative distribution function. We group firms based on their average lifetime probability of default and then study those in the top quartile with highest estimated default risk.

Table 1 presents the empirical results using both measures for closeness to exit. For all eight regressions, firms are on average procyclical, as given by the estimated coefficients for θ_1. Consider first the exit-based measure. As one would expect, firms that are in their final year before exit do less R&D on average relative to their lifetime — estimates of θ_3

[9]The market value of assets is unobservable and hence is estimated using a structural model iteratively (Black and Scholes, 1973) as

$$\begin{aligned} W_t &= V_t\Phi(d_1) - L_t e^{-r}\Phi(d_2) \\ \sigma_A &= \sigma(\log(V_t) - \log(V_{t-1})) \\ d_1 &= \frac{\log(V_t/L_t) + r + 1/2\sigma_A^2}{\sigma_A} \\ d_2 &= d_1 - \sigma_A \end{aligned}$$

where $\Phi(.)$ is the standard normal CDF, r is the real riskless rate, W_t is the market value of equity and $\sigma(.)$ is the standard deviation operator. Estimating this system gives us μ_A and σ_A for firms over their lifetime in the sample.

are stable around -0.3. Despite this, firms drifting towards exit undertake more state-contingent R&D — the interaction coefficient θ_2 is positive, statistically significant and large relative to the baseline output-growth coefficient. Note also that the inclusion of balance sheet controls lowers the estimate of θ_2 only modestly, indicating a force beyond simply relieving financing constraints. The distance-to-default measure points in the same direction. Since this measure applies to a broader set of firms — the highest-default-risk quartile — the estimated interaction coefficients are smaller than in the exit-based regressions. They are nonetheless large, positive and significant. Taken together, the results in Table 1 support selection over the firm lifecycle as a driver of R&D procyclicality. The pattern appears both immediately before observed exit and among firms persistently closer to the default margin.

Appendix B presents a variety of robustness checks and alternative exercises. Table A.1 shows that results are robust to including lags of the aggregate variable as well as interactions of lags with the sel_{it} variable. Table A.2 shows robustness with the same specification as (29), but with $\Delta out_{j(i)t} = \Delta out_t$ as economy-wide output growth, rather than industry-level growth. Table A.4 performs leave-one-out robustness over sectors. Leaving out manufacturing has the largest effect, given its high representation within Compustat; nevertheless, the interaction θ_2 remains positive and statistically significant. As a separate exercise motivated by the cross-sectional implication of Proposition 2, Table A.5 leverages a shorter balanced panel of firms and groups firms based on their average profitability, as a proxy for productivity. We find that firms in the lowest profitability quartile exhibit significantly greater R&D cyclicality. Table A.3 considers alternative outcome variables in specification (29). In particular, we also consider R&D intensity's correlation with the cycle, as well as the propensity for a firm to undertake R&D spending (i.e. the extensive margin). Firms about to exit the sample have positive interaction terms θ_2 for both measures. The results using the R&D propensity are particularly interesting since this allows us to include firms in the design, beyond the minority that report R&D.

5 Bringing the Model to the Data

In this section, we describe how the full model of Section 3 is mapped to the data. The calibration targets moments, which are chosen with two primary objectives in mind. The first is to discipline quantitatively the link between proximity to exit and R&D cyclicality documented in Section 4, particularly among firms with meaningful aggregate weight. The second objective is to then understand this mechanism's implications for aggregate growth.

Functional forms. Before discussing parameters, moments and identification, we specify the distributional assumptions used in the quantitative implementation. These as-

sumptions pertain to two objects that shape firm continuation values — the fixed cost distribution $F(\phi)$ and the distribution of innovation step sizes $H(\lambda)$. We assume fixed cost draws ϕ are lognormally distributed with location parameter $\mu_\phi \in \mathbb{R}$ and scale parameter $\sigma_\phi > 0$. Conditional on successful innovation, we assume the productivity step size λ is drawn from a Pareto distribution that implies minimum gross jump of one and shape parameter α_λ.

Parameters. The model contains the following 15 parameters

$$(\sigma, \beta, \eta, N, \delta, \alpha_\lambda, \mu_\phi, \sigma_\phi, \phi_e, \gamma, \gamma_e, \chi, \psi, \rho_Z, \sigma_Z).$$

We fix the preference parameters of the household externally to be standard values in the literature — with CRRA parameter $\sigma = 1$ and discount factor of $\beta = 0.96$ to give an annual interest rate of 4%. We set $\eta = 2.64$, a value broadly in line with empirical estimates of substitution across differentiated goods. For instance, Feenstra et al. (2018) report estimates that typically fall in the 2–4 range, depending on the level of aggregation and estimation method. We set exogenous exit in the baseline to zero $\delta = 0$ (this parameter will be used in later decomposition exercises). We normalize the detrended deterministic steady state wage to $\widehat{W} = 1$, with the sunk entry cost ϕ_e chosen so that the free entry condition holds. Given the resulting firm policies and equilibrium distribution, the labor endowment N is set equal to implied aggregate labor demand. We also normalize detrended output to $\widehat{Y} = 1$, which pins down the mass of firms in the stationary equilibrium.

The remaining nine parameters provide the quantitative margins used to match moments related to our mechanism. Since the calibration is solved jointly in general equilibrium, individual parameters should not be interpreted as being identified by individual moments. Instead, the parameter vector is chosen to match the full set of targeted moments using indirect inference. The calibration strategy involves a combination of matching deterministic steady state moments, as well as stochastic steady state moments through Monte Carlo simulation.[10] Panel A of Table 2 gives the parameter values, while Panel B gives the fit of the model moments to the data.

R&D cyclicality and selection mechanism. We calibrate to two moments relating to the responsiveness of R&D to aggregate shocks. First, we discipline the selection-cyclicality mechanism by asking whether firms accounting for the first 25 percent of sales, when cumulated from highest to lowest R&D procyclicality, are closer to the exit margin. Specifically, we run the regression Equation (1) estimated in Section 2, for each

[10] See Appendix C for details on the solution method.

Table 2: Calibration: parameter values and targeted moments.

Panel A: Parameter values

	Parameter	Value
Calibrated externally		
σ	CRRA parameter	1.000
β	Discount factor	0.960
η	Elasticity of substitution	2.639
Normalizations		
N	Labor supply	0.621
δ	Exogenous death	0.000
ϕ_e	Sunk entry cost	0.004
Calibrated internally		
α_λ	Shape parameter, step-size distribution	16.62
μ_ϕ	Location parameter, fixed-cost distribution	-2.565
σ_ϕ	Scale parameter, fixed-cost distribution	1.163
γ	Incumbent R&D efficiency	0.117
γ_e	Entrant R&D efficiency	2.027
χ	Productivity dependence of R&D costs	3.942
ψ	Curvature of R&D technology	1.262
ρ_Z	Persistence of aggregate TFP	0.950
σ_Z	Standard deviation of aggregate TFP	0.018

Panel B: Targeted moments

Target	Data	Model	Source
Deterministic steady state moments			
Growth rate	2.78%	2.78%	BEA
R&D to GDP	2.90%	2.88%	BEA
Exit rate	8.92%	8.92%	BDS
Entrant growth contribution	15.9%	15.9%	Pancost and Yeh (2022)
R&D intensity-size slope	-0.1169	-0.1183	Compustat
Stochastic steady state moments			
Average R&D cyclicality	0.3310	0.2935	Compustat
Exit premium, high-cyclicality sales group	17.1%	17.4%	Compustat
HP-filtered output autocorrelation	0.6073	0.6217	BEA
HP-filtered output standard deviation	0.0194	0.0192	BEA

Notes: Panel A reports the parameter values used in the quantitative model. The household preference parameters and elasticity of substitution are fixed externally. Labor supply and the sunk cost follow from normalizations. The number for ϕ_e is a percentage of steady state GDP (post multiplication by 100). The remaining parameters are chosen jointly using indirect inference to match the targeted moments in Panel B.

firm i in our Compustat sample. Recall this had the form

$$\Delta rd_{it} = a_i + \theta_{1,i}\Delta out_{j(i)t} + u_{it},$$

where $\theta_{1,i}$ measures the cyclicality of firm i's R&D growth with respect to output growth in its industry. We then rank firms from highest to lowest $\theta_{1,i}$, compute each firm's average lifetime real sales, then cumulate sales along this ranking until reaching 25% of total average lifetime sales of firms in the sample. Finally, we compare the average probability of default of firms (proxy for exit probability in the model) in this high-procyclicality sales group with that of all remaining firms.

Firms in the high-procyclicality group have a statistically significant premium in their probability of default of around 17%. Since this group accounts for 25% of sales by construction, matching this moment ensures that the calibrated mechanism applies to firms with a meaningful role in aggregate R&D dynamics. As a second moment, we match the pooled estimate of θ_1 in Specification (29) with our simulated firm sample, ensuring that the model reproduces the overall cyclicality of R&D. With intuition as suggested by (28), the R&D technology curvature ψ and parameters relating to fixed costs (μ_ϕ, σ_ϕ) are key margins in matching these moments. Higher values of ρ_Z also help match the exit risk premium in the high-procyclicality group, since persistence strengthens the continuation-value response to current aggregate conditions for firms close to exit.

As a validation test based on a moment not targeted in the parameterization, we use the simulated firm panel to estimate the distance-to-default interaction coefficient θ_2 from Specification (29). The interaction coefficient estimated from the simulated panel is statistically significant and reproduces around two thirds of the empirical magnitude reported in the first distance-to-default specification — the fifth numerical column of Table 1. Although untargeted, this result establishes that the calibrated cyclicality-selection mechanism also holds under the reverse sorting exercise. Whereas the calibration asks whether firms selected for high R&D cyclicality have greater default risk, this validation shows that firms selected for high default risk exhibit greater R&D cyclicality.

Firm aggregation. While the previous two moments discipline the responsiveness of R&D to aggregate shocks, we also target moments that determine the firm cross-section and level of R&D, through which these responses aggregate. We discipline the amount of selection in the economy by matching the overall firm exit rate from the BDS. We also target the relative contribution to growth of entrants to incumbents, as reported by Pancost and Yeh (2022). Parameters relating to the fixed cost distribution again serve as primary levers for matching the former moment, while the entrant-specific R&D productivity γ_e plays a central role in influencing the latter. We then also discipline the cross-sectional allocation of R&D by matching the negative relationship between firm

size and R&D intensity in Compustat referenced in Subfigure 1a in Section 2. Following Ignaszak and Sedláček (2026) and similar to Akcigit and Kerr (2018), we estimate

$$\log\left(1+\frac{rd_{it}}{sales_{it}}\right)=\nu\log(sales_{it})+\varepsilon_{it},$$

including sector and year fixed effects. We then target the estimated coefficient ν in the deterministic steady state of the model; parameter χ that enters the cost of R&D (7) through expression $\widehat{q}^{\chi}$ is a natural lever for matching this moment.

Aggregate innovation and business cycles. We discipline the scale and productivity of innovative activity by matching the R&D/GDP ratio of 2.9% and the average output growth rate of 2.78% reported by the BEA. A key object for targeting these moments is γ, the efficiency in converting resources to innovation for incumbents and α_λ which determines the average innovation step size. Since firm R&D decisions depend on expectations about the aggregate cycle, we also match the persistence and volatility of cyclical output. Specifically, we HP-filter output in both the model and the BEA data and match the persistence and volatility of the resulting cyclical component, with ρ_Z and σ_Z serving as the primary parameters to this end.

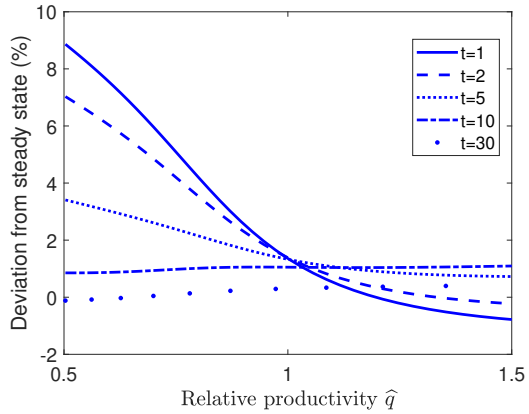
6 Quantitative analysis

We now turn from the mechanism to its aggregate quantitative implications. We proceed in three steps. First, we use the calibrated model to trace how the selection mechanism identified in Section 4 propagates through firm dynamics, aggregate output, and growth following an aggregate shock. Second, we decompose the resulting persistence into the contributions of endogenous growth and selection. Finally, we quantify how these channels affect the long-run output costs of a large recession.

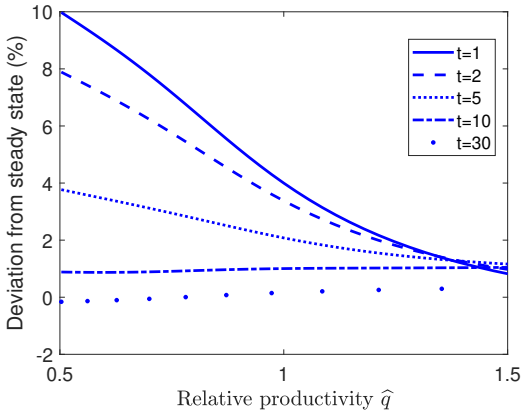
6.1 Selection, Innovation, and Aggregate Dynamics

Figure 3 traces the response of the calibrated economy to a one-standard-deviation positive aggregate TFP shock. Building on the qualitative validation of the analytical mechanism in Section 4, we now follow its transmission from heterogeneous firm-level innovation and survival responses to aggregate output and growth. Subfigure 3a shows the response of incumbent R&D choices, x, across the relative productivity distribution and over time. In the calibrated steady state, lower-productivity firms face a higher probability of exit. As predicted by the mechanism in the two period model of Section 4, these firms have a smaller innovation surplus and therefore respond more strongly to the positive shock to Z at $t=1$. For instance, a firm with $\widehat{q}=0.5$ increases x by around 9%, whereas a

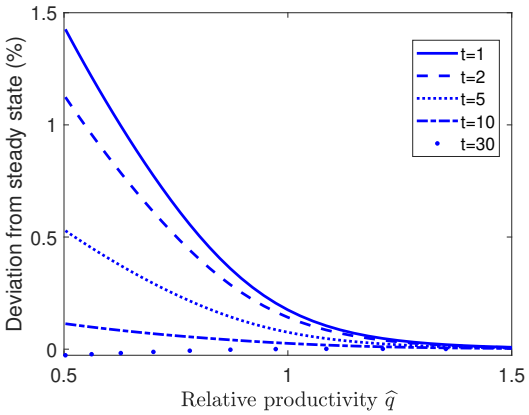
Figure 3: Baseline transition after an aggregate TFP shock.



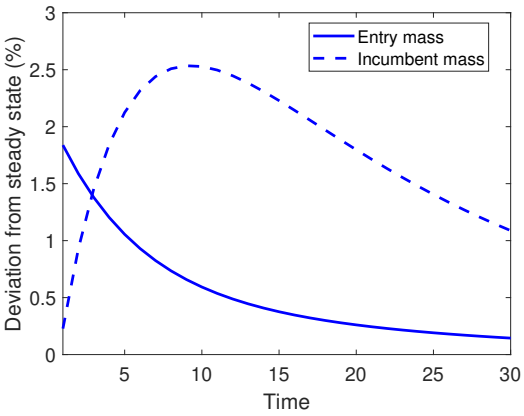
(a) Incumbent R&D policy



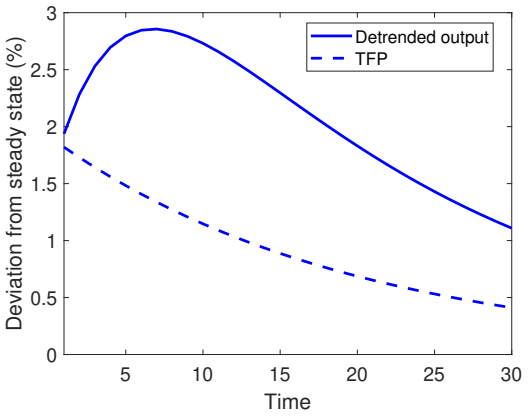
(b) Entrant R&D policy



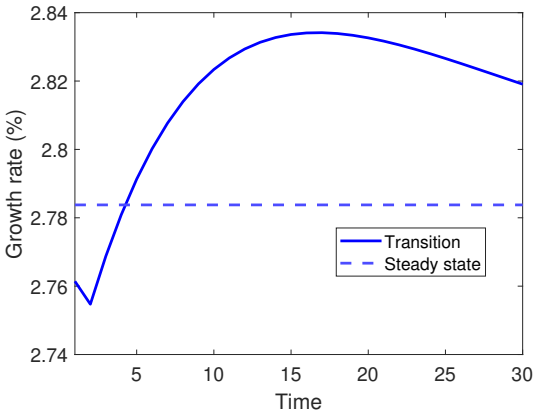
(c) Survival probabilities



(d) Firm masses



(e) TFP and detrended output



(f) Growth rate

Notes: The figure shows impulse responses along the baseline transition after an aggregate TFP shock. Subfigures 3a–3c plot responses as a function of relative productivity $\widehat{q}$ at selected horizons after the shock. The remainder give time profiles.

firm with $\widehat{q} = 1.5$ exhibits a countercyclical response of around 1%. The impact response therefore closely follows the prediction of the two-period model.

Further along the transition, however, Subfigure 3a shows a noticeable flattening of the cross-sectional response profile. The reason is that, in the infinite-horizon model, current R&D depends not only on next-period payoffs but on the entire sequence of future continuation values. Heuristically (and suppressing the common SDF movements), the linearized innovation response can be written as

$$\check{x}_t(\widehat{q}) \approx \frac{1}{\psi - 1} \sum_{h=1}^{\infty} K_h(\widehat{q}) \mathbb{E}_t \check{\pi}_{t+h},$$

where $K_h(\widehat{q})$ measures the sensitivity of current innovation incentives to payoffs h periods ahead. The two-period model of Section 4 contains only the one-period loading κ, whose cross-sectional variation is governed by firms' exposure to the exit margin. At longer horizons, by contrast, the relevant loadings incorporate all possible intermediate innovation and survival histories. As firms move across productivity states, their current position relative to the exit margin becomes progressively less informative about distant continuation values. The strong cross-sectional gradient predicted by the analytical mechanism is therefore most pronounced on impact and need not persist unchanged over the transition. In the calibrated economy it gradually flattens and is close to uniform by $t = 10$, before vanishing as the aggregate shock dissipates.[11]

Subfigure 3b shows a similar pattern for entrant R&D, but their responses are both larger and more persistent than those of incumbents. The simplified framework of Section 4 helps explain why. The entrant innovation response can be written as $\check{x}_e = 1/(\psi - 1)\mathbb{E}\{\check{M}' + \kappa_e \check{\pi}'\}$ where $\kappa_e = \bar{\pi}'_+ / \left(\bar{\pi}'_+ - \bar{c}'_+\right)$. The key difference is that an entrant becomes an incumbent only if innovation succeeds. Its return to R&D therefore depends on the absolute continuation value created by successful innovation. For an incumbent, instead, the return to innovation is the incremental value of success relative to failure, since the firm can continue operating under either outcome. The failure continuation value therefore partly offsets the response of the success value to aggregate conditions. No such offset is present for entrants. Aggregate cashflow movements consequently generate a larger elasticity of entrant R&D, particularly for low-$\widehat{q}$ entrants.

The same distinction also helps explain why entrant R&D remains more responsive during the early transition. For incumbents, the return to innovation is the difference between the successful and unsuccessful continuation values. As the direct effect of the shock on the exit margin fades, this difference becomes less sensitive to aggregate conditions and the cross-sectional profile of incumbent R&D responses flattens relatively quickly. Entrants, by contrast, value innovation through the level of the successful-entry value itself, with no corresponding failure value to offset its response. Their R&D response is therefore larger on impact and remains somewhat stronger at intermediate horizons.

[11] For example, the loading on payoffs two periods ahead contains terms such as $\bar{F}_+ \bar{x}_+ \bar{F}_{++}$ and $\bar{F}_+(1 - \bar{x}_+)\bar{F}_{+-}$, together with analogous terms along branches following unsuccessful innovation.

As aggregate conditions revert towards steady state, however, the entrant response profile also flattens, and by around ten years the incumbent and entrant profiles are quite similar.

Subfigure 3c shows the corresponding response of survival probabilities. The positive TFP shock generates an *anti-cleansing effect*: firms close to the exit margin become more likely to survive, while survival probabilities of more productive firms are essentially unchanged. As aggregate conditions revert towards steady state, this direct survival effect dissipates. Its consequences are nevertheless persistent, because the firms retained by the shock alter the cross-sectional distribution and the stock of active firms, both of which adjust only gradually.

Subfigure 3d shows that entry also rises sharply following the shock. Together with higher survival, this increases the mass of incumbents over the transition. Since the mass of firms is a state variable, these effects outlive the aggregate shock itself. The larger number of active varieties raises detrended output through the CES aggregator, so the persistent increase in firm mass translates into a persistent response of the stationary component of output, as shown in Subfigure 3e.

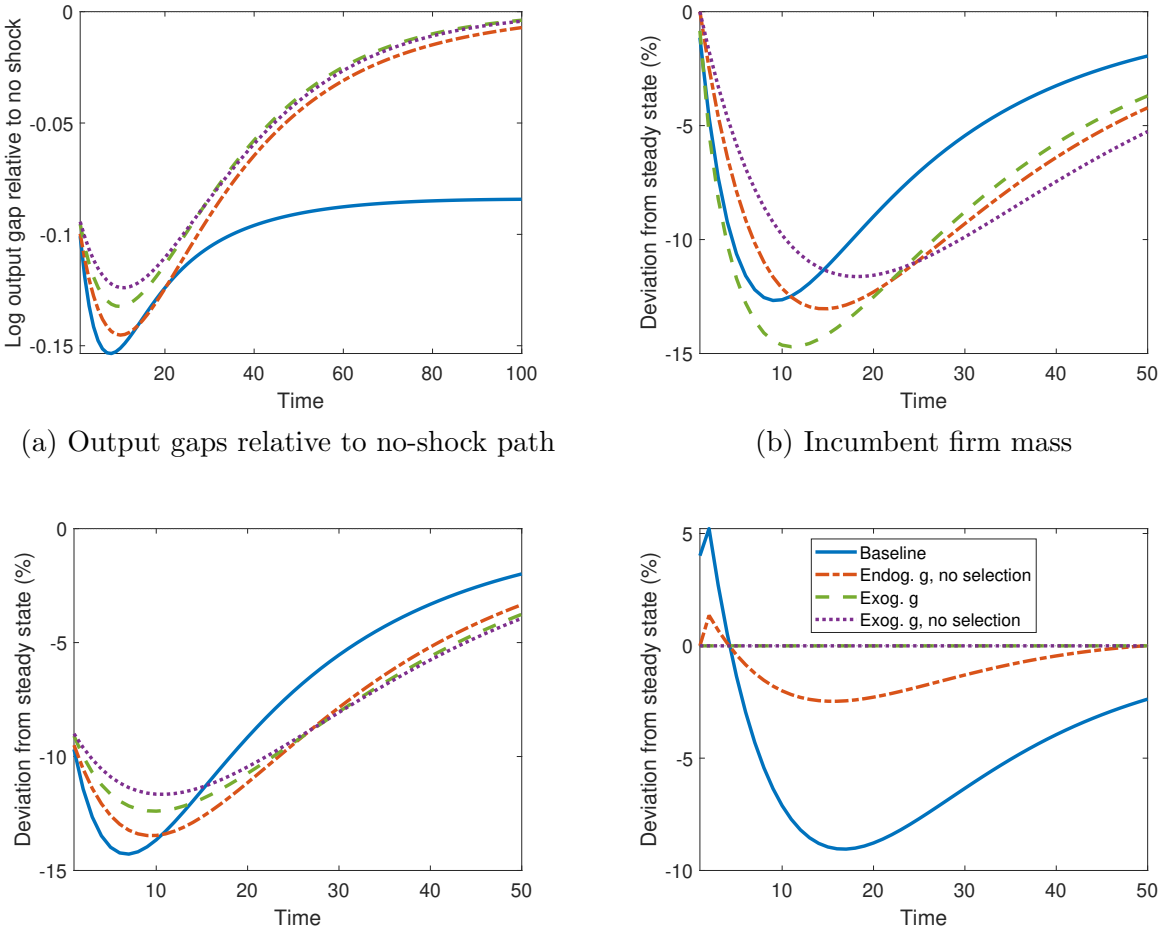
The response of growth in Subfigure 3f reflects two opposing forces. Initially, the anti-cleansing effect lowers productivity growth because the positive shock keeps relatively low-productivity firms in operation. The productivity aggregator Q, defined in Equation (8), measures average firm-level productivity and therefore neutralizes love-of-variety effects. The increased survival of lower-productivity firms consequently weighs directly on measured growth. After five years, however, the increase in R&D induced by the shock dominates this compositional effect. Productivity growth then rises above its steady-state rate and remains elevated even after the direct TFP shock has largely dissipated. This delayed growth response is a distinctive feature of endogenous innovation: in a model with exogenous growth, the shock can alter firm entry, exit, and composition, but cannot affect the subsequent path of productivity growth.

6.2 Decomposing the Long Shadow

We now use a set of counterfactual economies to separate the two forces that generate persistence in the baseline: endogenous growth and endogenous selection. The experiments can be viewed as a 2×2 decomposition according to whether (1) productivity growth is endogenous or exogenous and (2) whether endogenous firm selection is present or absent. This allows us to assess the propagation generated by firm demographics alone, the propagation generated by endogenous innovation in the absence of selection, and, most importantly, the interaction between selection and endogenous growth.

The first economy is the baseline calibration of Section 5, whose local dynamics were studied in Section 6.1. Starting from the deterministic steady state, we normalize $\widehat{Y} =$

Figure 4: Decomposing the transition following a large negative shock.



(a) Output gaps relative to no-shock path

(b) Incumbent firm mass

(c) Detrended output

(d) Endogenous TFP growth

Notes: The figure reports transition paths following the large negative shock. In each counterfactual, the initial shock is chosen to generate the same impact decline in output as in the baseline. Incumbent firm mass and detrended output are reported as percentage deviations from each economy's own steady state. Endogenous TFP growth is reported as the percentage deviation of the endogenous growth rate of average quality, $G_t = Q_t/Q_{t-1} - 1$, from its steady-state value; in the exogenous-growth economies, this response is zero by construction.

$Q = 1$ and subject the economy to a one-time negative aggregate TFP shock equal to five calibrated standard deviations. We then reconstruct the resulting path of output in levels, $Y = Q\widehat{Y}$, combining the transition in the stationary component $\widehat{Y}$ with the endogenous evolution of Q. The corresponding no-shock economy remains on its balanced-growth path and provides the benchmark against which output losses are measured.

We next construct an *exogenous growth* economy. Firms no longer undertake R&D and firm-level productivity instead follows an exogenous Markov process, while aggregate productivity grows at the baseline steady-state rate. Endogenous entry and exit are retained. We recalibrate the mean and variance of fixed operating costs, in addition to the persistence and variance of firm-level productivity to reproduce moments in the baseline parameterization. In particular — the baseline exit rate, expenditure on fixed costs/GDP, persistence of firm-level TFP and the gradient of the exit-productivity profile.[12] This economy therefore preserves the main firm-demographic and selection forces of the baseline while shutting down the endogenous-innovation channel. It provides an environment close in spirit to Clementi and Palazzo (2016), in which aggregate shocks propagate through entry, exit, and the evolving mass and distribution of firms.

We also construct an *endogenous growth, no selection* economy with symmetric R&D across firms. As shown in Section 4, removing fixed operating costs eliminates endogenous selection. Setting $\chi = \eta - 1$ then implies symmetric equilibrium R&D across the productivity distribution. We recalibrate the economy to match moments from the baseline deterministic steady state — the growth rate, median baseline incumbent innovation level x, exit rate (where exit is purely exogenous through δ) and γ_e such that entrants and incumbents undertake the same amount of R&D. We additionally match the correlation of growth in average R&D and output growth in the baseline stochastic Monte Carlo simulation, thereby ensuring approximately the same average R&D cyclicality.[13] The resulting economy is close to Barlevy (2007): long-run growth is endogenous and equilibrium innovation is symmetric across firms. Our framework generalizes this structure by allowing both entrants and incumbent firms to innovate. Relative to this benchmark, the baseline model features endogenous selection and the associated heterogeneity in R&D responses.

Finally, we construct the *exogenous growth, no selection* economy, which resembles the *exogenous growth* calibration but removes fixed operating costs and hence endogenous selection, with exit occurring only through the exogenous death rate δ. We set δ to match the baseline exit rate and recalibrate the dispersion of firm-level productivity so that the ratio of fixed-cost expenditure to GDP remains equal to that in the baseline, with fixed-cost expenditure in this case consisting entirely of sunk entry costs. Together, the four economies complete the 2×2 decomposition, allowing us to separately quantify the

[12] The details of the model and calibration can be found in Appendix D.

[13] For further details on the model and calibration, see Appendix E.

contributions of endogenous growth, selection and their interaction to the propagation of aggregate shocks. In each counterfactual, the magnitude of the aggregate shock is chosen to reproduce the same impact decline in output as in the baseline.

Panel 4a of Figure 4 reveals a sharp complementarity between endogenous growth and selection. Consider first the two exogenous-growth economies. Their output-gap paths are remarkably close. Endogenous selection therefore adds little propagation when productivity growth is exogenous. The result changes markedly when growth is endogenous. The baseline and the endogenous-growth, no selection economy initially display relatively similar output losses, but their paths subsequently diverge substantially. The baseline develops a much larger and more persistent scar. Selection therefore becomes quantitatively powerful precisely when it can affect firms' innovation decisions and, through them, the future path of productivity.

This result provides a useful comparison with Clementi and Palazzo (2016). Their mechanism emphasizes firm demographics: aggregate shocks change entry and exit, altering the mass and composition of firms, while the resulting cohorts evolve gradually as young firms grow through an exogenous productivity process. Our exogenous-growth economy retains endogenous entry and exit while shutting down endogenous innovation. In this environment, Panel 4a shows that introducing endogenous selection produces only a small additional long-run response.[14] Thus, in our exogenous-growth benchmark, selection and firm demographics alone are not a quantitatively powerful source of long-run propagation. By contrast, large propagation emerges once selection interacts with endogenous innovation, thereby affecting the future path of productivity.

Panel 4b helps explain the distinction. The contraction in the mass of incumbent firms is, if anything, stronger in the exogenous-growth economy than in the baseline. When innovation is endogenous, firms have an additional margin through which to respond to deteriorating aggregate conditions: they can reduce R&D. When that margin is absent, adjustment must instead occur more strongly through current production, entry and exit. The stronger demographic contraction in the exogenous-growth economy is consistent with this substitution across adjustment margins. Yet Panel 4a shows that this larger movement in firm mass does not translate into a larger long-run output scar. The demographic effects gradually unwind as the mass and distribution of firms recover.

Panel 4c reinforces this conclusion. Detrended output falls sharply following the crisis, reflecting the response of production, entry, exit and the cross-sectional distribution of active firms. These effects are persistent because the mass and distribution of firms are state variables and adjust only gradually. They do not, however, account for the much larger long-run output loss in the baseline. At longer horizons, detrended output in the

[14] Clementi and Palazzo (2016) also feature firm-level physical capital subject to non-convex adjustment costs, which provides an additional slow-moving margin of adjustment. Our exogenous-growth benchmark abstracts from this friction.

baseline recovers relatively strongly even as the overall output gap in Panel 4a remains substantially larger. The source of the long shadow must therefore lie in the growth component of output.

Panel 4d identifies this source by showing the response of endogenous productivity growth. Immediately after the crisis, productivity growth rises slightly in the baseline. This is the familiar "cleansing effect" of recessions (Caballero and Hammour, 1994): tighter selection disproportionately eliminates low-productivity firms and temporarily raises average productivity growth. But the gain is short-lived. As the contraction in innovation feeds into future productivity improvements, growth turns downward and remains below its steady-state rate for an extended period. In the endogenous-growth, no selection economy, the corresponding decline in productivity growth is much smaller. There is still a long shadow in the absence of selection, as temporarily lower growth cumulates into a persistent productivity loss, but this effect is substantially less severe than in the baseline, as shown in 4a. The temporary cleansing gain is therefore eventually overwhelmed by the cumulative effect of "missing innovation".

The comparison with the no-selection economy clarifies why endogenous growth alone is not sufficient to generate the baseline's long shadow. In the symmetric-R&D environment — close to the endogenous-growth business-cycle framework of Barlevy (2007) — a recession reduces aggregate innovation and thereby affects future productivity. Although this economy is calibrated to reproduce the same average R&D cyclicality as the baseline, it does not reproduce the underlying heterogeneity in firm-level responses. In the baseline, firms closer to the exit margin have substantially more cyclical R&D, while endogenous selection changes the composition of firms following the shock. These forces jointly generate a more persistent contraction in aggregate innovation and productivity growth. Selection therefore amplifies the long-run consequences of the initial downturn even when average R&D cyclicality is approximately matched.

Importantly, firms close to exit also tend to be smaller and, as in the data, smaller firms are relatively more R&D intensive (Akcigit and Kerr, 2018). A negative shock therefore disproportionately pushes these firms toward exit, shifting the composition of active firms toward less R&D-intensive firms and reducing aggregate innovation intensity. Moreover, the heterogeneous R&D response reinforces this initial selection effect over time. By cutting R&D more strongly, firms close to the exit margin are less likely to innovate successfully, remain closer to the exit margin and become more likely to exit in subsequent periods. The resulting additional selection therefore further reduces aggregate innovation intensity and TFP growth. Selection therefore affects propagation both directly, through an initial cleansing effect of small firms, and indirectly, through the feedback between heterogeneous R&D responses and subsequent selection. This feedback further depresses aggregate innovation and TFP growth, generating substantially greater long-run propagation than in the symmetric-R&D economy.

Table 3: Output losses following a crisis.

Model	Year 10	Year 20	Year 30	Year 50
Baseline	15.11	12.40	10.56	9.07
Endogenous growth, no selection	14.52	12.43	9.18	4.48
Exogenous growth	13.24	11.36	8.38	3.83
Exogenous growth, no selection	12.39	11.04	8.41	4.02

Notes: Log output losses at the indicated horizon relative to each model's no-shock path. Numbers come from the time paths of output presented in Figure 4a.

Taken together, the four experiments identify a strong interaction between firm selection and endogenous growth. Selection without endogenous growth generates persistent movements in firm demographics, but contributes relatively little to the long-run output scar. Endogenous growth without selection allows a temporary recession to affect future productivity, but the resulting scar is substantially smaller than in the baseline. It is their interaction that generates the pronounced long shadow: selection amplifies the cyclical response of innovation, while endogenous growth converts the resulting missing innovation into a cumulative productivity loss. Selection generates heterogeneous R&D cyclicality at the firm level; endogenous growth turns that heterogeneity into aggregate persistence.

The quantitative importance of this amplification speaks to a broader debate on the extent to which severe downturns leave persistent scars on economic activity. A large empirical literature finds that major recessions are often followed by persistent output losses (e.g. Cerra and Saxena, 2008). The COVID-19 pandemic renewed interest in the possibility that unusually large temporary disruptions can have effects lasting well beyond the initial contraction (Jordà et al., 2022). One potential source of such persistence is endogenous technical change: recessions reduce innovation and technology adoption, lowering future productivity and thereby turning temporary shocks into persistent output losses (Anzoategui et al., 2019; Benigno and Fornaro, 2018).

To put numbers to the importance of the endogenous growth-selection mechanism, we interpret the large negative shock studied in Figure 4 as a COVID-sized crisis; Table 3 reports the magnitudes of the output path comparisons. Recall that the one-time negative TFP shock was five calibrated standard deviations — approximately a 9 log-point decline in TFP and generating an approximately 10% fall in baseline real GDP on impact. This broadly matches the magnitude of the initial contraction associated with the pandemic. We do not interpret this experiment as a structural model of the pandemic itself; rather, COVID-19 provides a useful benchmark for the magnitude of an unusually severe aggregate contraction. We then use this same experiment to compare the resulting numerical output losses across the four economies introduced above.

Allowing for endogenous growth substantially increases the cost of the crisis, but the

magnitude of this effect depends critically on selection. Recall that the two exogenous-growth economies isolate the firm-demographics channel (e.g. Clementi and Palazzo, 2016). Consistent with the transition paths above, the quantitative differences between these two economies are small. At the 50-year horizon, for example, output losses are 3.83 log points with selection and 4.02 without it. Selection by itself is therefore not a quantitatively important source of long-run scarring in this experiment.

The picture changes markedly once growth is endogenous. At the 50-year horizon, the baseline economy remains 9.07 log points below its no-shock path, compared with only 4.48 log points in the endogenous-growth economy without selection. Thus, removing selection cuts the 50-year output scar by roughly one half. Selection therefore becomes quantitatively important precisely when firms' endogenous innovation decisions determine the subsequent path of productivity.

The results reveal a strong complementarity between endogenous growth and selection. Selection has only modest and, at some horizons, non-monotonic effects on the output gap when productivity growth is exogenous. Its consequences become large when innovation determines future productivity. A severe downturn pushes more firms towards the exit margin, where R&D is especially sensitive to aggregate conditions. The resulting contraction in innovation lowers subsequent growth and leaves a much larger long-run scar. The same selection mechanism that generates heterogeneous R&D cyclicality at the firm level therefore materially amplifies the aggregate output costs of a large recession.

7 Conclusion

The cyclical response of R&D differs substantially across firms. Motivated by this new empirical fact, we develop a theory that identifies endogenous selection as a source of heterogeneous innovation responses to aggregate shocks. Our key analytical result is that firm heterogeneity alone is not sufficient to generate heterogeneous innovation elasticities: firms can differ in productivity, size, profits, and steady-state R&D while responding proportionally to aggregate fluctuations. Endogenous selection breaks this neutrality by making the value of innovation particularly sensitive to aggregate conditions for firms close to the exit margin. We test this mechanism directly in the data and find that firms closer to exit exhibit significantly more procyclical R&D, even after controlling for balance-sheet characteristics.

The quantitative analysis reveals a strong interaction between selection and endogenous growth. When productivity growth is exogenous, selection and changes in firm demographics generate relatively little additional long-run propagation. When growth is endogenous, however, selection affects firms' innovation decisions and thereby the future path of productivity. Our counterfactual exercises show that removing selection while retaining endogenous growth reduces the output scar of a severe recession by roughly one

half at a 50-year horizon.

These findings highlight a dimension of business-cycle propagation that is largely absent from existing endogenous-growth models. The long-run consequences of a recession depend not only on the aggregate decline in innovation, but also on its distribution across firms. By shaping which firms adjust innovation most strongly, endogenous selection can turn temporary aggregate shocks into substantially larger and more persistent productivity losses.

These findings also open up new questions for policy. In particular, understanding how stabilization and innovation policies interact with firm selection, and whether they can mitigate the long-run productivity losses generated by recessions, is a promising avenue for future research.

References

Acemoglu, Daron, Ufuk Akcigit, Nicholas Bloom, and William Kerr, "Innovation, Reallocation and Growth," *American Economic Review*, 2018, *108* (11), 3450–3491.

Aghion, Philippe, Marios Angeletos, Abhijit Banerjee, and Kalina Manova, "Volatility and Growth: Credit Constraints and Productivity-Enhancing Investment," *Journal of Monetary Economics*, 2010, *57*, 24–67.

_ , Nicolas Berman, Laurent Eymard, Philippe Askenazy, and Gilbert Cette, "Credit Constraints and the Cyclicality of R&D Investment: Evidence from France," *Journal of the European Economic Association*, 2012, *10*, 1001–1024.

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

Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez, "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence," *American Economic Journal: Macroeconomics*, 2019, *11* (3), 67–110.

Ates, Sina T. and Felipe E. Saffie, "Fewer but Better: Sudden Stops, Firm Entry, and Financial Selection," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 304–356.

Ball, Lawrence, "Long-Term Damage from the Great Recession in OECD Countries," *NBER Working Paper*, 2014, (20185).

Barlevy, Gadi, "On the Cyclicality of Research and Development," *American Economic Review*, 2007, *97* (4), 1131–1164.

Benigno, Gianluca and Luca Fornaro, "Stagnation Traps," *Review of Economic Studies*, 2018, *85*, 1425–1470.

Bertolotti, Fabio and Andrea Lanteri, "Capital replacement and innovation dynamics," Technical Report, CEPR Discussion Papers 2024.

Bianchi, Francesco, Howard Kung, and Gonzalo Morales, "Growth, slowdowns, and recoveries," *Journal of Monetary Economics*, 2019.

Black, Fischer and Myron Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 1973, *81* (3), 637–654.

Blanchard, Olivier, Eugenio Cerutti, and Lawrence Summers, "Inflation and Activity—Two Explorations and Their Monetary Policy Implications," *NBER Working Paper*, 2015, (21726).

Caballero, Ricardo J. and Mohamad L. Hammour, "The Cleansing Effect of Recessions," *The American Economic Review*, 1994, *84* (5), 1350–1368.

Cerra, Valerie and Sweta Chaman Saxena, "Growth Dynamics: The Myth of Economic Recovery," *American Economic Review*, 2008, *98* (1), 439–457.

_ , Antonio Fatas, and Sweta Saxena, "Hysteresis and Business Cycles," *Journal of Economic Literature*, 2023, *2023*, 181–215.

Clementi, Gianluca and Bernardino Palazzo, "Entry, Exit, Firm Dynamics and Aggregate Fluctuations," *American Economic Journal: Macroeconomics*, 2016, *8* (3), 1–41.

Comin, Diego and Mark Gertler, "Medium-Term Business Cycles," *American Economic Review*, 2006, *96* (3), 532–551.

Cooper, Russel and John Haltiwanger, "On the Nature of Capital Adjustment Costs," *Review of Economic Studies*, 2006, *73* (3), 611–633.

Corbae, Dean and Pablo D'Erasmo, "Reorganization or liquidation: Bankruptcy choice and firm dynamics," *The Review of Economic Studies*, 2021, *88* (5), 2239–2274.

Crouzet, Nicolas and Neil R. Mehrotra, "Small and Large Firms over the Business Cycle," *American Economic Review*, 2020, *110* (11), 3549–3601.

Fatas, Antonio, "Do Business Cycles Cast Long Shadows? Short-run Persistence and Economic Growth," *Journal of Economic Growth*, 2000, *5*, 147–162.

Feenstra, Robert C., Philip Luck, Maurice Obstfeld, and Katheryn N. Russ, "In Search of the Armington Elasticity," *The Review of Economics and Statistics*, 2018, *100* (1), 135–150.

Fornaro, Luca and Martin Wolf, "The Scars of Supply Shocks: Implications for Monetary Policy," *Journal of Monetary Economics*, 2023, *140*, S18–S36.

Fort, Teresa, John Haltiwanger, Ron S. Jarmin, and Javier Miranda, "How firms respond to business cycles: The role of firm age and firm size.," *IMF Economic Review*, 2013, *3* (61), 520–559.

Gertler, Mark and Simon Gilchrist, "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *Quarterly Journal of Economics*, 1994, *109* (2), 309–340.

Griliches, Zvi, "Patent Statistics as Economic Indicators," *Journal of Economic Literature*, 1990, *28*.

Hopenhayn, Hugo A., "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 1992, *60* (5), 1127–1150.

Ignaszak, Marek and Petr Sedláček, "Customer acquisition, business dynamism and aggregate growth," *Review of Economic Studies*, 2026, p. rdag060.

Jordà, Òscar, Sanjay R. Singh, and Alan M. Taylor, "Longer-Run Economic Consequences of Pandemics," *The Review of Economics and Statistics*, 2022, *104* (1), 166–175.

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

Krusell, Per and Anthony A Smith Jr, "Income and wealth heterogeneity in the macroeconomy," *Journal of political Economy*, 1998, *106* (5), 867–896.

Ottonello, Pablo and Thomas Winberry, "Financial Heterogeneity and the Investment Channel of Monetary Policy," *Econometrica*, 2020, *88* (6), 2473–2502.

Ouyang, Min, "On the Cyclicality of R&D," *Review of Economics and Statistics*, 2011, *93*, 542–533.

Pancost, N Aaron and Chen Yeh, "Decomposing Aggregate Productivity," Technical Report 2022.

Queralto, Albert, "A Model of Slow Recoveries from Financial Crises," *Journal of Monetary Economics*, 2020, *114*, 1–25.

Schmitz, Tom, "Endogenous Growth, Firm Heterogeneity and the Long-run Impact of Financial Crises," *European Economic Review*, 2021, *132* (2).

Sedláček, Petr, "Creative destruction and uncertainty," *Journal of the European Economic Association*, 2020, *18* (4), 1814–1843.

Spencer, Adam Hal, "Policy effects of international taxation on firm dynamics and capital structure," *The Review of Economic Studies*, 2022, *89* (4), 2149–2200.

Winberry, Thomas, "Lumpy Investment, Business Cycles, and Stimulus Policy," *American Economic Review*, 2021, *111* (1), 364–396.

A Simplified Framework Derivations

Here we provide the steps used to obtain the first-order approximation in Equation (28). Recall from Equation (26) that the time-zero value function in the two-period model can be written as

$$v = \max_{s} \left\{ \pi + \mathbb{E}\widehat{M}' \left[x\tilde{v}'_{+} + (1-x)\tilde{v}'_{-} \right] \right\},$$

where, as in Section 4.1,

$$\begin{aligned}\widetilde{v}'_{\pm} &= F'_{\pm}\left(\pi'_{\pm} - c'_{\pm}\right),\\ F'_{\pm} &\equiv F(\pi'_{\pm}),\\ c'_{\pm} &\equiv \frac{1}{F'_{\pm}}\int_0^{\pi'_{\pm}} \phi\, dF(\phi).\end{aligned}$$

The R&D technology is $s = \gamma \widehat{q}^{\chi} x^{\psi}$. Hence the first-order condition for innovation is

$$\psi\gamma\widehat{q}^{\chi}x^{\psi-1} = \mathbb{E}\widehat{M}'\left[\widetilde{v}'_{+} - \widetilde{v}'_{-}\right],$$

which is the condition underlying Equation (27).

We first linearize terminal equity values. Since

$$\widetilde{v}'(\pi') = \int_0^{\pi'} (\pi' - \phi)\, dF(\phi),$$

we have

$$\frac{\partial \widetilde{v}'(\pi')}{\partial \pi'} = F(\pi').$$

Therefore, for each terminal innovation outcome,

$$d\widetilde{v}'_{\pm} = \bar{F}'_{\pm}\, d\pi'_{\pm} = \bar{F}'_{\pm}\bar{\pi}'_{\pm}\check{\pi}'_{\pm}.$$

Equivalently,

$$\bar{\widetilde{v}}'_{\pm}\check{\widetilde{v}}'_{\pm} = \bar{F}'_{\pm}\bar{\pi}'_{\pm}\check{\pi}'_{\pm},$$

where bars denote deterministic steady-state values and checks denote percentage deviations from steady state.

We now linearize the first-order condition for x. Write

$$\begin{aligned}x &= \bar{x}\exp(\check{x}),\\ \widetilde{v}'_{\pm} &= \bar{\widetilde{v}}'_{\pm}\exp(\check{\widetilde{v}}'_{\pm}),\\ \widehat{M}' &= \bar{\widehat{M}}\exp(\check{M}').\end{aligned}$$

Substituting into the first-order condition and using $\exp(z) \simeq 1 + z$ gives the steady-state condition

$$\psi\gamma\widehat{q}^{\chi}\bar{x}^{\psi-1} = \bar{\widehat{M}}\left(\bar{\widetilde{v}}'_{+} - \bar{\widetilde{v}}'_{-}\right)$$

and the first-order condition

$$\psi\gamma\widehat{q}^{\chi}\bar{x}^{\psi-1}(\psi-1)\check{x} = \widehat{\bar{M}}E\Bigg[\left(\check{\bar{v}}'_{+}-\check{\bar{v}}'_{-}\right)\check{M}' + \check{\bar{v}}'_{+}\check{\bar{v}}'_{+}-\check{\bar{v}}'_{-}\check{\bar{v}}'_{-}\Bigg].$$

Dividing the second expression by the steady-state condition yields

$$\check{x} = \frac{1}{\psi-1}\mathbb{E}\left[\check{M}' + \frac{\check{\bar{v}}'_{+}\check{\bar{v}}'_{+}-\check{\bar{v}}'_{-}\check{\bar{v}}'_{-}}{\check{\bar{v}}'_{+}-\check{\bar{v}}'_{-}}\right].$$

Substituting the linearized terminal equity values and the definitions of $\check{v}'_{\pm}$ gives

$$\check{x} = \frac{1}{\psi-1}\mathbb{E}\left[\check{M}' + \frac{\bar{F}'_{+}\bar{\pi}'_{+}\check{\pi}'_{+}-\bar{F}'_{-}\bar{\pi}'_{-}\check{\pi}'_{-}}{\bar{F}'_{+}\left(\bar{\pi}'_{+}-\bar{c}'_{+}\right)-\bar{F}'_{-}\left(\bar{\pi}'_{-}-\bar{c}'_{-}\right)}\right].$$

It remains to use the CES cashflow expression. Equation (25), evaluated in the terminal period together with the law of motion for relative productivity, implies that successful and unsuccessful innovators differ in the level of terminal cashflow but have the same percentage exposure to aggregate variables:

$$\check{\pi}'_{+} = \check{\pi}'_{-} = \check{\pi}',$$

where

$$\check{\pi}' = -(\eta-1)\check{W}' + (\eta-1)\check{Z}' + \check{Y}' - (\eta-1)\frac{\bar{G}}{1+\bar{G}}\check{G}'.$$

We can therefore write

$$\check{x} = \frac{1}{\psi-1}\mathbb{E}\left[\check{M}' + \kappa\check{\pi}'\right],$$

where

$$\kappa = \frac{\bar{F}'_{+}\bar{\pi}'_{+}-\bar{F}'_{-}\bar{\pi}'_{-}}{\bar{F}'_{+}\left(\bar{\pi}'_{+}-\bar{c}'_{+}\right)-\bar{F}'_{-}\left(\bar{\pi}'_{-}-\bar{c}'_{-}\right)}.$$

A.1 Proof of Proposition 2

It remains to characterize how the amplification varies with productivity. Let

$$\tau = (\eta-1)\log(1+\bar{\lambda})$$

denote the log proportional increase in terminal cashflow generated by successful innovation. Then

$$\bar{\pi}'_+(\widehat{q}) = e^{\tau}\bar{\pi}'_-(\widehat{q}).$$

By Equation (25) and the terminal law of motion for relative productivity, $\bar{\pi}'_-(\widehat{q})$ is strictly increasing in $\widehat{q}$. Define

$$\zeta = \log \bar{\pi}'_-(\widehat{q}), \qquad g(\ell) \equiv F(e^{\ell}).$$

Using the value-gain representation in Section 4.1, evaluated at $a = 1$ for a firm with productivity $\widehat{q}$, and the change of variables $z = e^{\zeta}e^{t}$,

$$\begin{aligned}\Delta(1) &= e^{\zeta}\int_0^{\tau} e^{t}g(\zeta + t)\,dt \\ &\equiv e^{\zeta}H(\zeta),\end{aligned}$$

where

$$H(\zeta) = \int_0^{\tau} e^{t}g(\zeta + t)\,dt.$$

Since a common proportional change in terminal cashflows is equivalent to a change in ζ, the loading in Equation (28) can be written as

$$\kappa(\widehat{q}) = 1 + \frac{d\log H(\zeta)}{d\zeta}.$$

Suppose that $g(\ell) = F(e^{\ell})$ is log-concave. Then

$$(\zeta, t) \mapsto e^{t}g(\zeta + t)$$

is jointly log-concave, because

$$\log\left(e^{t}g(\zeta + t)\right) = t + \log g(\zeta + t)$$

is concave in (ζ, t). By Prekopa's theorem, which preserves log-concavity under integration, $H(\zeta)$ is log-concave. Hence

$$\frac{d^2\log H(\zeta)}{d\zeta^2} \leqslant 0,$$

and consequently

$$\frac{d}{d\zeta}\left[1+\frac{d\log H(\zeta)}{d\zeta}\right]\leqslant 0.$$

Since $\bar{\pi}'_{-}(\widehat{q})=e^{\zeta}$ is increasing in both ζ and $\widehat{q}$, $\kappa(\widehat{q})$ is non-increasing in relative productivity.

For the lognormal fixed-cost distribution used in the quantitative analysis,

$$F(\phi)=\Phi\left(\frac{\log\phi-\mu_{\phi}}{\sigma_{\phi}}\right),$$

so that

$$g(\ell)=F(e^{\ell})=\Phi\left(\frac{\ell-\mu_{\phi}}{\sigma_{\phi}}\right).$$

The Gaussian CDF is log-concave, and log-concavity is preserved by affine transformations. The sufficient condition above is therefore satisfied by the lognormal fixed-cost distribution. This proves Proposition 2.

ONLINE APPENDIX

B Robustness for Empirical Support

Table A.1: Selection mechanism regressions robustness with lagged output growth.

	Exit-based measure				Distance-to-default measure			
Output growth θ_1	0.299***	0.311***	0.529**	0.553**	0.265***	0.285***	0.401***	0.412***
	(0.082)	(0.083)	(0.204)	(0.196)	(0.085)	(0.093)	(0.125)	(0.126)
Interaction θ_2	1.926***	2.012***	1.917***	2.002***	0.259**	0.324***	0.332***	0.397***
	(0.296)	(0.512)	(0.342)	(0.548)	(0.090)	(0.092)	(0.091)	(0.088)
Close to exit θ_3	−0.317***	−0.302***	−0.352***	−0.338***				
	(0.008)	(0.022)	(0.010)	(0.018)				
Lagged output growth θ_4	0.346***	0.355**	0.536**	0.579**	0.268**	0.240**	0.420***	0.438***
	(0.108)	(0.128)	(0.206)	(0.205)	(0.099)	(0.096)	(0.131)	(0.127)
Lagged interaction θ_5	0.423**	0.540	0.190	0.330	0.107	0.175*	0.226**	0.300**
	(0.194)	(0.443)	(0.188)	(0.408)	(0.065)	(0.092)	(0.096)	(0.113)
Time FE	Y	Y	N	N	Y	Y	N	N
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Controls	Y	N	Y	N	Y	N	Y	N
Observations	64,590	67,689	64,590	67,689	55,226	57,037	55,226	57,037

Notes: The dependent variable in all specifications is real R&D growth. Columns 1–4 use the exit-based measure and columns 5–8 use the distance-to-default measure. All specifications include contemporaneous and lagged output growth, each interacted with the relevant measure of closeness to exit. Standard errors are in parentheses and clustered at the industry level. Stars *, ** and *** denote significance at the 10%, 5% and 1% levels. The estimated specification augments Equation (29) to include lagged output growth:

$$\Delta rd_{it} = a_i + \theta_1 \Delta out_{j(i)t} + \theta_2 (sel_{it} \times \Delta out_{j(i)t}) + \theta_3 sel_{it} + \theta_4 \Delta out_{j(i),t-1} + \theta_5 (sel_{it} \times \Delta out_{j(i),t-1}) + u_{it}.$$

Table A.2: Selection mechanism regressions robustness using aggregate output growth.

	Exit-based measure		Distance-to-default measure	
Output growth θ_1		2.300***		1.749***
		(0.521)		(0.273)
Interaction θ_2	4.153***	3.843***	0.936***	1.654***
	(1.296)	(1.243)	(0.206)	(0.273)
Close to exit θ_3	−0.408***	−0.430***		
	(0.055)	(0.049)		
Time FE	Y	N	Y	N
Firm FE	Y	Y	Y	Y
Controls	Y	Y	Y	Y
Observations	64,590	64,590	55,226	55,226

Notes: The dependent variable in all specifications is real R&D growth using Equation (29). This table repeats the baseline regressions using aggregate GDP growth rather than industry output growth. Columns 1–2 use the exit-based measure and columns 3–4 use the distance-to-default measure. All specifications include firm fixed effects and balance-sheet controls. In columns with time fixed effects, the main effect of aggregate output growth is absorbed by the time effects and is therefore not reported. Standard errors are in parentheses and clustered at the industry level. Stars *, ** and *** denote significance at the 10%, 5% and 1% levels.

Table A.3: Selection mechanism regressions robustness with alternative dependent variables.

	R&D intensity		R&D propensity	
Output growth θ_1	−0.456	−5.161***	−1.103	−0.695
	(1.438)	(1.227)	(0.860)	(0.420)
Interaction θ_2	9.799***	6.412***	1.181***	1.107***
	(2.494)	(2.029)	(0.308)	(0.251)
Close to exit θ_3	0.284	0.183	−0.118***	−0.129***
	(0.186)	(0.136)	(0.039)	(0.037)
Time FE	Y	N	Y	N
Firm FE	Y	Y	N	N
Controls	Y	Y	Y	Y
Observations	64,590	64,590	182,000	182,000

Notes: This table reports exit-based robustness specifications using alternative dependent variables. Columns 1–2 use R&D intensity and columns 3–4 use an indicator for positive R&D expenditure. All specifications include balance-sheet controls. Columns 1 and 3 include time fixed effects; columns 2 and 4 omit time fixed effects. Firm fixed effects are omitted in the R&D participation specifications because R&D participation exhibits very limited within-firm variation. Standard errors are in parentheses and clustered at the industry level. Stars *, ** and *** denote significance at the 10%, 5% and 1% levels. Specifications replace the left-side variable in Equation (29).

Table A.4: Leave-one-out sector robustness.

Omitted sector	Interaction θ_2	Observations
Mining, quarrying, oil and gas extraction	2.027***	64,012
	(0.377)	
Construction	1.993***	64,344
	(0.351)	
Manufacturing	3.862*	15,224
	(1.923)	
Wholesale	1.982***	63,852
	(0.347)	
Retail	2.054***	64,318
	(0.363)	
Transport and warehousing	1.996***	64,545
	(0.350)	
Information	1.740***	56,805
	(0.160)	
Finance and insurance	2.006***	63,957
	(0.350)	
Real estate, rental and leasing	2.010***	63,920
	(0.357)	
Professional, scientific and technical services	2.054***	61,367
	(0.372)	
Administrative and support,	1.905***	64,108
waste management and remediation services	(0.315)	
Health care and social assistance	1.998***	64,212
	(0.350)	
Arts, entertainment and recreation	1.999***	64,509
	(0.351)	
Accommodation and food services	1.995***	64,497
	(0.350)	

Notes: The dependent variable is real R&D growth. The table reports the interaction coefficient from the baseline exit-based specification after omitting one sector at a time. All specifications include firm fixed effects, time fixed effects and balance-sheet controls. Standard errors are in parentheses and clustered at the industry level. Stars *, ** and *** denote significance at the 10%, 5% and 1% levels. Specifications follow the form of Equation (29).

Table A.5: Firm cyclicality regressions grouped by profit ratios

Output growth θ_1	-0.257	-0.228	0.306*	0.313**
	(0.551)	(0.528)	(0.128)	(0.124)
Interaction θ_2	0.437**	0.384**	0.376***	0.326***
	(0.123)	(0.124)	(0.071)	(0.070)
Time FE	Y	Y	N	N
Firm FE	Y	Y	Y	Y
Controls	Y	N	Y	N
Observations	7,337	7,337	7,337	7,337

Notes: The dependent variable in all specifications is real R&D growth. The sample is a balanced panel covering 1989–2018. The interaction is between industry output growth and an indicator for firms in the bottom quartile of lifetime average profitability within their industry. All columns use the same common sample. Standard errors are in parentheses and clustered at the industry level. Stars *, ** and *** denote significance at the 10%, 5% and 1% levels.

$$\Delta rd_{it} = a_i + \theta_1 \Delta out_{j(i)t} + \theta_2 \left(m_i \times \Delta out_{j(i)t}\right) + u_{it},$$

where m_i is an indicator for a firm being in the lowest profitability quartile, a_i is a firm-level fixed effect and u_{it} is a residual. The coefficient of interest here is θ_2, which gives an indication of the incremental cyclicality of firms that are the least profitable in the balanced sample.

C Computational Algorithms

In this appendix, we detail the solution of the deterministic steady state and transition after a shock, in turn.

C.1 Deterministic Steady State

The following describes the computation of the calibrated steady state. Given it is deterministic, we omit time subscripts. Note here that $Z = 1$ and the aggregate state Γ is constant — we therefore drop it in the notation that follows.

1. Create a grid for relative productivity levels $\widehat{q}$.

2. Fix certain aggregates: this reduces the iterative complexity of the problem. Specifically, we fix the detrended wage $\widehat{W} = 1$ as well as detrended aggregate output $\widehat{Y} = 1$. To make these aggregates consistent with our equilibrium definition, we treat the sunk cost of entry ϕ_e as a free parameter and find the implied firm mass Ω. Recall also that the aggregate price index is the numéraire in our model, and as such we set $P = 1$.

3. Solve for the static choices $p(\widehat{q})$, $n(\widehat{q})$ and $\widehat{y}(\widehat{q})$ of incumbent intermediate-goods firms, consistent with Step 2 and the expressions in Equation (16).

4. Solve the incumbent intermediate goods firms' Bellman equation (13) using value function iteration. This uses the static controls from Step 3 as inputs. This step yields their policy functions for R&D $x(\widehat{q})$ and $s(\widehat{q})$, the fixed-cost cutoff $\widetilde{\phi}(\widehat{q})$, expected fixed costs, exit probabilities and value functions $v(\widehat{q})$ and $\widetilde{v}(\widehat{q})$. Note that we extrapolate the value function beyond the top grid point, so that firms close to the maximum $\widehat{q}$ grid point still have an incentive to innovate, as in the model.

5. Solve the entrant's problem using the value function $\widetilde{v}(\widehat{q})$ found in Step 4 as an input and the expressions given in Equations (17) and (18). This yields entrants' optimal R&D choices $x_e(\widehat{q}_e)$ and $s_e(\widehat{q}_e)$, as well as entrant value conditional on prototype productivity, $v_e(\widehat{q}_e)$.

6. Find the sunk cost of entry consistent with the free-entry condition, the object $v_e(\widehat{q}_e)$ found in Step 5 and the assumption that $\widehat{W} = 1$ from Step 2:

$$\phi_e = \mathbb{E}_e \left[v_e(\widehat{q}_e) \right].$$

7. Find the stationary endogenous cross-section of firms over relative productivity. Do this by first assuming a unit measure of firms. We can write Equation (22) in

matrix notation as

$$\boldsymbol{\omega}_t = \mathcal{P}_t \boldsymbol{\omega}_{t-1} + \mathcal{P}_{e,t} \Omega_{e,t} \boldsymbol{\omega}_{e,t}, \tag{30}$$

where $\boldsymbol{\omega}_t$ is the vector of incumbent firm masses over the relative-productivity grid, $\mathcal{P}_t$ is the endogenous incumbent transition matrix implied by the equilibrium objects found in Steps 4–5, $\mathcal{P}_{e,t}$ is the entrant transition matrix, $\Omega_{e,t}$ is the mass of potential entrants and $\boldsymbol{\omega}_{e,t}$ is the distribution of prototype relative productivities among successful entrants. The stationary distribution can then be written as

$$\widetilde{\boldsymbol{\omega}} = \widetilde{\Omega}_e \left(I - \mathcal{P}\right)^{-1} \mathcal{P}_e \boldsymbol{\omega}_e,$$

where I is the identity matrix and $\widetilde{\Omega}_e$ initially normalizes the total firm mass to unity.

8. Find variable averages across all incumbent intermediate-goods firms that are consistent with the distribution from Step 7 and the optimal controls from Steps 3–5.

9. Find the mass of entrants consistent with the normalization $\widehat{Y} = 1$ from Step 2, using average detrended output from Step 8 and the stationary distribution from Step 7, leveraging the linearity of the latter. This also implies the cross-sectional measure of incumbent firms over the relative-productivity grid.

10. Compute the aggregate variables consistent with Step 9. Find aggregate labor demand from intermediate-goods firms, L. Set the household's inelastic labor supply equal to this object, $N = L$, and treat this as a parameter in the transition simulations.

11. Find the aggregate growth rate using Equation (24).

C.2 Transition After TFP Shock

To solve for the dynamic equilibrium, we build on Sedláček (2020), extending the methodology to allow for endogenous growth. In particular, we use a first-order perturbation around the stationary steady state. The approximated solutions take the form

$$\begin{aligned}
\widehat{\boldsymbol{h}}_{t+1} &= \overline{\boldsymbol{h}} + \Theta_1 \left(\widehat{\boldsymbol{h}}_t - \overline{\boldsymbol{h}}\right), \\
\widehat{\boldsymbol{d}}_{t+1} &= \overline{\boldsymbol{d}} + \Theta_2 \left(\widehat{\boldsymbol{h}}_t - \overline{\boldsymbol{h}}\right),
\end{aligned}$$

where Θ_1 and Θ_2 contain the coefficients obtained from the approximation, $\widehat{\boldsymbol{h}}_t$ contains the predetermined state variables and $\widehat{\boldsymbol{d}}_t$ contains the non-predetermined variables. The perturbation is carried out jointly over the aggregate state and the full cross-sectional

distribution of firms.[15]

D Exogenous Growth Model Details

The firm production function is changed to

$$y_t(j) = T_t Z_t \zeta_t(j) n_t(j)$$

where $\zeta_t(j)$ is firm-specific exogenous productivity process and T_t is a common aggregate growth trend. The firm-specific (dropping j subscripts) productivity process evolves according to

$$\log(\zeta_t) = \rho_\zeta \log(\zeta_{t-1}) + \epsilon_{\zeta,t}$$

for $0 < \rho_\zeta < 1$ and $\epsilon_{\zeta,t} \sim N(0, \sigma_\zeta^2)$. For shorthand, we denote this exogenous process as $\mathcal{K}(\zeta_t|\zeta_{t-1})$. New entrants draw their initial productivity from the ergodic distribution of this process, denoted $\mathcal{K}_e(\zeta)$. The aggregate growth component evolves according to

$$T_t = (1 + \bar{g})T_{t-1}$$

where $\bar{g}$ is the exogenous growth rate. We then de-trend variables with respect to T_t. All remaining elements of the household and final-goods blocks are unchanged. In particular, household preferences and the budget constraint are given by Equations (2) and (3), with the stochastic discount factor defined from Equation (10). Final-goods producers continue to use the CES aggregator in Equation (4), implying the variety-level demand function and price index in Equations (11) and (12), respectively, with the final good serving as the numéraire. Aggregate TFP Z_t continues to follow the process in Equation (6). Intermediate-goods firms remain monopolistic competitors, face the same exogenous death probability δ and draw i.i.d. fixed operating costs from $F(\phi)$, while potential entrants continue to pay the sunk entry cost ϕ_e. The Bellman equation for an incumbent can be written as

$$\begin{aligned}
v(\zeta, \Gamma) &= \max_{p,n} \pi(\zeta, \Gamma) + (1-\delta)\mathbb{E}\widehat{M}'\widetilde{v}(\zeta', \Gamma'), \\
\pi(\zeta, \Gamma) &= p\,\widehat{y} - \widehat{W}n, \\
\widetilde{v}(\zeta, \Gamma) &= \int_\phi \max\{v(\zeta, \Gamma) - \phi, 0\}\, dF(\phi),
\end{aligned}$$

[15] This approach has the advantage of being computationally convenient relative to global methods for models with aggregate shocks, such as Krusell and Smith (1998) and the MIT shock solution method used in Spencer (2022).

subject to

$$
\begin{aligned}
p(\zeta,\Gamma) &= \frac{\eta}{\eta-1}\frac{\widehat{W}}{Z\zeta},\\
\widehat{y}(\zeta,\Gamma) &= p(\zeta,\Gamma)^{-\eta}\widehat{Y},\\
n(\zeta,\Gamma) &= \frac{\widehat{y}(\zeta,\Gamma)}{Z\zeta}\\
\widetilde{\phi}(\zeta,\Gamma) &= v(\zeta,\Gamma).
\end{aligned}
$$

In the exogenous-growth economy, level profits and firm equity values inherit the common trend T_t. Throughout this section, $\pi(\zeta,\Gamma)$, $v(\zeta,\Gamma)$ and $\widetilde{v}(\zeta,\Gamma)$ denote these objects after division by T_t (we omit the hats in the interest of notational brevity). Similarly, fixed operating costs and the sunk entry cost are likewise expressed in units of T_t. The law of motion for the firm cross section is

$$
\omega'(\zeta',Z') = (1-\delta)\int_{\zeta}\int^{\widetilde{\phi}(\zeta',\Gamma')} \mathcal{K}(\zeta'|\zeta)\omega(\zeta,Z)\,dF(\phi)\,d\zeta + \Omega_e\int^{\widetilde{\phi}(\zeta',\Gamma')} \mathcal{K}_e(\zeta')\,dF(\phi).
$$

The first term captures incumbent firms that survive the exogenous death shock, transition from productivity ζ to ζ' according to $\mathcal{K}(\zeta'|\zeta)$ and remain active after drawing their next-period fixed operating cost. The second term captures new entrants, whose initial productivity is drawn from $\mathcal{K}_e(\zeta')$, and which survive their initial fixed-cost draw. Entrants' conditional value is

$$
v_e(\zeta_e,\Gamma) = \mathbb{E}\widehat{M}'\widetilde{v}(\zeta_e,\Gamma'),
$$

and the ex-ante value to entry is

$$
\widetilde{v}_e(\Gamma) = -\phi_e + \mathbb{E}_e\left[v_e(\zeta_e,\Gamma)\right],
$$

where the expectation $\mathbb{E}_e$ is taken over entrants' initial productivity draws from $\mathcal{K}_e(\zeta_e)$. Free entry requires $\widetilde{v}_e(\Gamma)=0$. Equilibrium is then defined similarly to the list in Section 3.2. For the purpose of the simulation experiments in Section 6.2, note that since each intermediate producer's output scales with T_t, aggregate output does as well. Using $\widehat{y}_t(j) = y_t(j)/T_t$ and $\widehat{Y}_t = Y_t/T_t$ therefore gives

$$
\begin{aligned}
\widehat{Y}_t &= \left(\int_{j\in\Omega_t}\widehat{y}_t(j)^{\frac{\eta-1}{\eta}}\,dj\right)^{\frac{\eta}{\eta-1}}\\
Y_t &= T_t\widehat{Y}_t.
\end{aligned}
$$

Table A.6: Exogenous growth calibration: parameter values and targeted moments.

Panel A: Parameter values

Parameter			Value
Calibrated externally		Selection	No selection
σ	CRRA parameter	1.000	1.000
β	Discount factor	0.960	0.960
η	Elasticity of substitution	2.639	2.639
ρ_Z	Persistence of aggregate TFP	0.950	0.950
σ_Z	Standard deviation of aggregate TFP	0.018	0.018
ρ_ζ	Persistence of firm productivity	0.975	0.975
Normalizations			
N	Labor supply	0.621	0.621
ϕ_e	Sunk entry cost	229.3	148.1
Calibrated internally			
σ_ζ	Std. dev. firm productivity	0.193	0.063
μ_ϕ	Location fixed-cost distribution	-0.322	
σ_ϕ	Scale fixed-cost distribution	1.055	
δ	Exogenous death rate	0.000	0.089

Panel B: Targeted moments

Target	Data	Selection	No selection	Source
Growth rate	2.78%	2.79%	2.79%	BEA
Exit rate	8.92%	8.85%	8.92%	BDS
Fixed-cost expenditure/output	0.233	0.257	0.258	Baseline
Corr. productivity and survival	0.772	0.699		Baseline

Notes: Panel A reports the parameter values used in the two exogenous-growth counterfactuals. The household preference parameters and elasticity of substitution are fixed externally. Labor supply and the sunk cost follow from normalizations. The remaining parameters are chosen to match the targeted moments in Panel B.

E Endogenous Growth, No Selection Calibration

Table A.7: Counterfactual calibration: parameter values and targeted moments.

Panel A: Parameter values

	Parameter	Value
Calibrated externally		
σ	CRRA parameter	1.000
β	Discount factor	0.960
η	Elasticity of substitution	2.639
ρ_Z	Persistence of aggregate TFP	0.950
σ_Z	Standard deviation of aggregate TFP	0.018
Normalizations		
N	Labor supply	0.621
ϕ_e	Sunk entry cost	0.840
Calibrated internally		
α_λ	Shape parameter, step-size distribution	8.517
ψ	Curvature of R&D technology	1.060
γ	Incumbent R&D efficiency	0.186
γ_e	Entrant R&D efficiency	1.087
δ	Exogenous death rate	0.089

Panel B: Targeted moments

Target	Data	Model	Source
Growth rate	2.78%	2.79%	BEA
Exit rate	8.92%	8.92%	BDS
Incumbent x	0.199	0.201	Baseline
Entrant x_E	0.199	0.166	N/A
Corr. growth of ave. R&D with output	0.940	0.880	Baseline

Notes: Panel A reports the parameter values used in the quantitative model. The household preference parameters and elasticity of substitution are fixed externally. Labor supply and the sunk cost follow from normalizations. The number for ϕ_e is a percentage of steady state GDP (post multiplication by 100). The remaining parameters are chosen jointly using indirect inference to match the targeted moments in Panel B.