

The Distribution of Innovators and Medium-Term Business Cycles*

Lixing Wang[†]

Donghai Zhang[‡]

September, 2026

Abstract

Firm-level R&D changes are usually small, with occasional large increases that are procyclical and predict patent quantity and quality. We develop a medium-term business-cycle model with heterogeneous innovators who accumulate raw ideas and face fixed project-launch costs. Past shocks reshape their raw-idea distribution, making subsequent shock propagation history dependent. Low volatility leaves more innovators close to launching, so subsequent contractions postpone more projects and cause larger losses. In the calibrated model, identical Great Recession shocks generate long-run cumulative TFP losses of 4.7% starting from the pre-recession distribution, versus 3.3% in a counterfactual using the average distribution. The model traces this vulnerability to the preceding Great Moderation, generating the Great Twins hypothesis. Cross-country evidence supports this prediction; firm-level R&D data support the mechanism involving the pre-recession concentration of innovators near the launch threshold. The central bank therefore faces a trade-off between reducing business-cycle volatility and limiting the severity of subsequent recessions.

Keywords: Medium-run Business Cycles, Heterogeneous Innovators, Great Moderation, Great Recession

JEL Classification: E32, E52, O31

*We thank Diego Comin, John Fernald, Pablo Guerrón-Quintana, Pablo Ottonello, Petr Sedláček, Vincent Sterk, Daniel Xu, as well as the seminar/conference participants at Boston University, Northwestern University, Boston College, Federal Reserve Board, Cleveland Fed, University of New South Wales, Collegio Carlo Alberto, Bank of Portugal, University of Bonn, Singapore Management University, National University of Singapore, Fudan University, SUFE, Jinan University (IESR), Academia Sinica, Peking University, NBER EASE 2024, Virtual East Asian Macroeconomics Seminar, and Southern Workshop in Macroeconomics for their suggestions and discussions. Donghai Zhang thanks the Singapore Ministry of Education for financial support under the SSRTG grant SSRC2023-SSRTG-026. Lixing Wang thanks the DFG for financial support as part of the research training group RTG 2281 'The Macroeconomics of Inequality'.

[†]National University of Singapore, wang-lixing@outlook.com

[‡]National University of Singapore, donghai.d.zhang@gmail.com

1 Introduction

Temporary contractions can have lasting effects when they depress innovation and technology adoption (Comin and Gertler, 2006; Comin et al., 2009). Existing models generally explain these effects through changes in aggregate R&D spending and the creation and adoption of new technologies. Innovation, however, is organized around discrete projects. Its response to shocks therefore also depends on how close different innovators are to starting a new project: a hidden cross-sectional distribution that aggregate innovation measures do not capture. Past shocks and stabilization policy can reshape this distribution, which in turn governs the response to future shocks. This raises a question: what are the implications of these distributional dynamics for business-cycle propagation, stabilization policy, and growth?

We first document three patterns in firm-level R&D. First, annual changes in reported R&D expenditure are usually small, with occasional large increases. Second, the share of firms making these large increases is procyclical. Third, these increases predict subsequent patent counts and several measures of patent importance and value. Together, these facts show that R&D adjustment is episodic, with a procyclical extensive margin that predicts subsequent innovation outcomes.

We then develop a business-cycle framework with endogenous innovation and heterogeneous innovators, focusing on the distribution of innovators by how close they are to launching a project. Each innovator accumulates raw ideas, which are ideas not yet assigned to an active R&D project.¹ Starting a project requires a fixed setup cost. An innovator therefore launches when its stock of raw ideas reaches a threshold that varies with aggregate conditions. Innovators below the threshold carry their stocks forward and continue to accumulate ideas. At launch, all accumulated raw ideas are assigned to the new project, and the innovator's raw-idea stock resets to zero. Project building and adoption then convert launched projects into new technologies.

The model's key aggregate state is the distribution of raw-idea stocks across innovators. This distribution records both how many innovators are close to starting a project and how many raw ideas they hold. Two economies with similar levels of aggregate innovation can have different masses of innovators near project launch and therefore respond differently to the same aggregate shock.

¹Throughout the paper, an innovator's raw-idea stock can be interpreted as either the quantity of raw ideas or the quality of an idea; the two interpretations are isomorphic.

Project launches make the distribution path dependent. A favorable shock causes some innovators to launch and reset their raw-idea stocks, while an adverse shock causes others to postpone launch. When aggregate conditions reverse, earlier launches cannot be undone. The distribution therefore retains the effects of past shocks. We refer to this history dependence as *distributional hysteresis*. It provides a cross-sectional mechanism through which past shocks shape the propagation of future shocks (Cerra et al., 2023).

The key result is that distributional hysteresis is asymmetric. When project launch is sufficiently lumpy, more innovators lie immediately below the launch threshold than immediately above it. A favorable shock shifts the threshold leftward, inducing launches and resets. An adverse shock that shifts the threshold equally far rightward postpones fewer launches (Lemma 1). Repeated threshold movements change the *shape* of the distribution of innovators, as formalized in Proposition 2. High aggregate volatility therefore leaves fewer innovators near the launch threshold, whereas periods of lower aggregate volatility, which we call moderation, allow this concentration to build.

The inherited distribution governs how contractions propagate. When an adverse shock raises the threshold, the decline in new project launches depends on how many innovators are displaced from the launch margin and how many raw ideas they hold. After a period of moderation, more innovators are close to launching. The same contraction then postpones more projects and leaves more raw ideas outside active R&D. Project building and adoption turn this launch shortfall into a gradual and persistent decline in technology, productivity, and output.

Stabilization policy therefore changes the distribution on which future shocks operate. A central bank that places greater emphasis on stabilization, represented in the model by a more aggressive Taylor rule, reduces business-cycle fluctuations and thereby changes the history of project launches. Given the asymmetry described above, greater stabilization leaves more innovators near the launch threshold. The distribution of innovators is thus endogenous to macroeconomic policy. This mechanism generates the *Great Twins* hypothesis: the Great Moderation contributed to the severity of the Great Recession.

To quantify the mechanism, we calibrate the model to firm-level patterns in lumpy R&D and fit it to U.S. GDP growth. Holding fixed the shocks that triggered the Great Recession, the estimated pre-recession distribution produces a long-run cumulative TFP loss of 4.7%, compared with 3.3% in a counterfactual starting from the model's time-averaged stationary distribution. Increasing volatility during 2002–2007 by 20% reduces the cumulative TFP loss caused by the same recessionary shocks by 1.5 percentage points. More

generally, the model predicts that greater emphasis on stabilization lowers unconditional volatility but produces a distribution under which a severe contraction that monetary policy cannot offset causes a larger and more persistent loss. The central bank therefore faces a trade-off between reducing business-cycle volatility and limiting downside tail risk.

Policy-induced changes in the distribution also affect average growth. Greater moderation shifts the distribution toward larger raw-idea stocks and leaves these ideas outside active R&D projects for longer. Because unused raw ideas can become obsolete, this longer delay causes more of them to lose value before project launch. Greater stabilization can therefore reduce fluctuations at the cost of lower average growth.

Finally, we document the Great Twins pattern in the data and provide evidence supporting its distributional mechanism. Countries that experienced greater moderation before 2008 also suffered deeper recessions during the Great Recession. This evidence is correlational, but the relationship remains positive and statistically significant after controlling for cross-country differences in real interest rates and pre-recession credit expansion. More importantly, the positive relationship appears only among R&D-intensive economies, consistent with the model's innovation mechanism. The Great Twins pattern also appears across post-World War II recessions.

At the micro level, we provide evidence for the mechanism operating through the pre-recession distribution of innovators. We use the distribution of time elapsed since firms' last R&D spike as a proxy for the distribution of raw-idea stocks. Consistent with the model, the probability of an R&D spike increases with elapsed time, supporting the use of this proxy. The 2008 distribution of elapsed time closely matches its model counterpart. This match supports the model's mechanism involving the concentration of innovators near the launch threshold before the Great Recession.

Literature Review Our paper contributes to the literature on medium-term business cycles pioneered by [Comin and Gertler \(2006\)](#). They document joint high- and medium-frequency movements in output, R&D, and technological change and develop a unified model in which cyclical R&D and technology adoption propagate nontechnology shocks through endogenous productivity. Earlier work linked business-cycle fluctuations to endogenous technological change and long-run growth ([Stadler, 1990](#); [Fatas, 2000](#)). [Barlevy \(2007\)](#) explains how dynamic externalities in innovation can make R&D procyclical by inducing firms to focus on short-term returns. Related work links innovation and technology adoption to expectations of future growth and asset prices ([Comin et al., 2009](#); [Kung and Schmid, 2015](#)). Other studies examine how weak demand depresses innovation and

technology adoption, with persistent effects at the zero lower bound (Anzoategui et al., 2019; Benigno and Fornaro, 2018), and how financial constraints affect innovation and growth (Aghion et al., 2010; Guerron-Quintana and Jinnai, 2019; Queralto, 2020; Bianchi et al., 2019; Schmitz, 2021; Ottonello and Winberry, 2026). Across these models, shocks alter innovation and the future path of technology. We instead ask how past fluctuations and stabilization policy reshape the cross-sectional distribution governing later innovation responses. To address this question, we extend this class of models to heterogeneous innovators who accumulate raw ideas and face fixed project-launch costs.

Firm entry and composition can also amplify and prolong the aggregate effects of shocks (Bilbiie et al., 2012; Sedláček and Sterk, 2017; Sedláček, 2020). In Bilbiie et al. (2012), sunk entry costs propagate shocks through slow adjustment in the number of producers. Sedláček and Sterk (2017) show that business-cycle conditions at entry change startups' growth potential, while Sedláček (2020) show that recessions create missing cohorts whose employment effects persist as they age. Persistent growth types also shape firm dynamics and the aggregate effects of firm-level frictions (Sterk et al., 2021). In contemporaneous work, Benedetti-Fasil et al. (2025) document that more productive firms' innovation is less procyclical and model selection as the source of this heterogeneity. Coupled with firm entry and exit, heterogeneous R&D cyclicalities amplify the propagation of TFP shocks. Together, this literature makes firm composition an inherited aggregate state. Our mechanism instead operates within continuing innovators: launch-and-reset decisions make their raw-idea distribution an inherited state shaped by past shocks and stabilization policy.

The lumpy-investment literature studies whether fixed adjustment costs matter for aggregate dynamics and how the distribution of capital gaps shapes responses to shocks (Caballero et al., 1995; Caballero and Engel, 1999; Thomas, 2002; Khan and Thomas, 2008; Koby and Wolf, 2020). Quantitative work shows that lumpy adjustment can generate procyclical and history-dependent responses (Bachmann et al., 2013; Winberry, 2021). Baley and Blanco (2021) show that the transitional dynamics of aggregate capital following productivity shocks are structurally linked to two steady-state moments: the dispersion of capital-to-productivity ratios and their covariance with the time elapsed since the last adjustment. The closest antecedent to our history-dependence result is Lee (2026), in which a surge in lumpy investments by large firms leaves many of them recently adjusted and makes subsequent TFP-induced recessions more severe. In our paper, raw ideas accumulate toward a one-sided project-launch threshold and reset at launch. When project

launch is sufficiently lumpy, the resulting distributional hysteresis is asymmetric. In contemporaneous work, [Berlingieri et al. \(2026\)](#) document that firms create products in bursts and show how the distribution of burst sizes shapes firm-size concentration and growth. Their focus is the cross-sectional concentration of product creation. We instead study the cyclical timing of lumpy innovation: the frequency of large R&D adjustments is procyclical in the data, and aggregate shocks endogenously shift the project-launch margin in the model. This cyclical extensive margin makes the distribution of innovators an inherited state governing medium-term business-cycle propagation.

The paper proceeds as follows. Section 2 documents episodic R&D adjustment and the procyclicality of large increases. Section 3 develops the project-launch mechanism, establishes asymmetric distributional hysteresis, and shows how aggregate fluctuations reshape the distribution of innovators. Section 4 embeds this mechanism in a general equilibrium business-cycle model. Section 5 calibrates the model and recovers historical shocks and distributions using U.S. GDP growth. Section 6 quantifies how the pre-recession distribution amplified the Great Recession and traces this amplification to the preceding Great Moderation. It also examines the implications for stabilization policy and average growth. Section 7 presents cross-country evidence for the Great Twins and firm-level evidence for the distributional mechanism, and shows that large R&D increases predict subsequent patenting.


2 Motivating Evidence

This section documents two features of firm-level R&D adjustment that motivate the model. First, annual changes in reported R&D expenditure are usually small, with occasional large increases. Second, the share of firms making these large increases is procyclical.

Data Description Our empirical analysis uses annual firm-level panel data from S&P Capital IQ’s Compustat database for 1968–2020. We focus on U.S. firms and exclude firms in the financial and utility sectors. We define an innovating firm as one that reports positive R&D expenditure at least once during the sample period. For the time-series analysis, we retain firms with at least five observations. We also exclude firm-year observations in which acquisitions exceed 5% of total assets.² The resulting sample contains 7,214 inno-

²Acquisitions may incorporate the R&D departments of acquired companies, causing abrupt changes in reported R&D expenditure. Our results are robust to retaining observations excluded by this acquisition

Figure 1: Episodic R&D Adjustment



Notes: This figure plots the firm-level distribution of percent changes in R&D expenses. To accommodate observations with zero initial R&D expenses, we calculate the percent change as $\frac{R\&D_t - R\&D_{t-1}}{0.5 \times (R\&D_t + R\&D_{t-1})}$. Excess kurtosis measures the deviation of the fourth moment from that of a normal distribution. The excess tail ratio compares the empirical 80th-to-95th quantile ratio with its counterpart under a normal distribution. P-values are obtained by bootstrapping.

vating firms, representing 20.4% of the Compustat universe. The panel is unbalanced, with an average of 18.9 observations per firm.

The sample covers a similar set of industries as the full Compustat population but places greater weight on R&D-intensive manufacturing firms. Table B.1 in the Online Appendix reports the industry composition. Our primary variable is annual R&D expense. It includes salaries and other costs associated with developing, designing, and testing new products or improving existing products. It also includes amortized software costs and other capitalized R&D assets. We deflate R&D expenses using the GDP deflator. Throughout the paper, R&D expenses refer to real values.

Episodic R&D Adjustment Figure 1 plots the distribution of annual percent changes in firm-level R&D expenses. Most observations cluster around small adjustments, while the distribution has a pronounced right tail of large increases. The excess kurtosis is 5.15, rejecting normality with a p-value of 0.00. The excess tail ratio, defined as the percentage deviation of the empirical 80th-to-95th quantile ratio from that of a normal distribution, is 70%. These patterns show that firm-level R&D adjustment is episodic: most annual changes are small, but firms occasionally make large increases.

The episodic pattern is stable across several alternative samples and measures. Online screen.

Table 1: Cyclicalities of Large R&D Adjustments

	Baseline	Alternative Spike Thresholds			
Spike Threshold	40%	20%	60%	80%	100%
Corr. with GDP Growth	0.54	0.45	0.54	0.50	0.35

Notes: The R&D spike rate is the annual share of innovating firms whose R&D expenditure increases by more than the indicated threshold. Changes below \$10,000 and observations in which acquisitions exceed 5% of total assets are excluded.

Appendix B.2 reports the supporting evidence and describes these robustness checks in detail. Briefly, the pattern remains when we normalize R&D by sales, restrict the sample to firms whose sales changed by less than 10% over the preceding three years, focus on firms with large R&D expenditures, split the sample by financial leverage, and exclude entry and exit observations. Annual changes in patent counts also display a pronounced right tail. R&D tax incentives are unlikely to generate this pattern. The U.S. credit for increasing research activities applies to qualified expenditures above a historically determined base amount, with a minimum tied to current qualified expenditures. Its design therefore rewards sustained research spending rather than sharp year-to-year shifts in spending.

Cyclicalities of Large R&D Adjustments We next examine whether large R&D adjustments vary with aggregate conditions. We define an R&D spike as an annual increase in R&D expenditure above a chosen threshold. We exclude changes below \$10,000 and firm-year observations in which acquisitions exceed 5% of total assets. The R&D spike rate is the annual share of innovating firms making a spike adjustment.

Table 1 reports the cyclicalities of the spike rate across alternative thresholds. Under our baseline choice of 40%, the correlation between the spike rate and real GDP growth is 0.54. Alternative thresholds from 20% to 100% deliver similar results: all correlations are positive and range from 0.35 to 0.54. Online Appendix Figure B.4 plots the pooled-threshold spike rate alongside real GDP growth, while Figure B.3 shows similar procyclicality among firms with small and large R&D expenditures. This finding contributes to the empirical literature on R&D cyclicalities (Fatas, 2000; Barlevy, 2007; Broda and Weinstein, 2010; Ouyang, 2011) by highlighting the importance of the extensive margin of R&D adjustment.

These patterns motivate a model in which large R&D adjustments arise from discrete project launches whose timing responds to aggregate conditions.

3 Heterogeneous Innovators: Key Mechanisms

We study heterogeneous innovators who accumulate raw ideas and pay a fixed cost to launch a new R&D project. A launch assigns the innovator's accumulated raw ideas to the project and resets the raw-idea stock. The distribution of these stocks across innovators is therefore an aggregate state shaped by past project launches.

The innovation block builds on [Romer \(1990\)](#). There is a continuum of innovators with a total measure of unity. Each innovator creates new varieties of intermediate inputs and earns profits by selling the associated intellectual property rights. [Figure 4](#) summarizes the three phases within each period. In Phase I, innovators receive raw ideas and decide whether to launch a new R&D project. In Phase II, ideas assigned to active projects—which we call R&D ideas—undergo random adoption. This phase requires skilled labor, whose cost accounts for most R&D expenditure in the model. In Phase III, adopted ideas become new varieties of intermediate inputs whose intellectual property rights are sold to intermediate-input producers.

This section isolates the Phase I project-launch decision. The partial-equilibrium analysis shows how launch decisions reshape the distribution of raw-idea stocks. [Section 4](#) completes Phases II and III and embeds the mechanism in a general equilibrium model with business-cycle shocks.

3.1 Innovators' Decisions in Phase I

Each innovator enters period t with a stock of raw ideas $s_{i,t}$. The innovator decides whether to pay a fixed setup cost c and launch a new project. We refer to this project-launch decision as adjustment in the formal problem below.

When an innovator does not launch, the stock evolves according to

$$s_{i,t+1} = \phi s_{i,t} + \eta_{i,t+1} \tag{3.1}$$

where $1 - \phi \in (0, 1)$ is the obsolescence rate of unused ideas and $\eta_{i,t+1}$ is the inflow of new ideas. For the analytical results in this section, idea inflows are exogenous and independently and identically distributed as $N(\mu, \tau^2)$ draws conditional on being nonnegative, where $\mu \in \mathbb{R}$ and $\tau > 0$ are the location and standard deviation before truncation. [Section 4](#) endogenizes idea creation and embeds the mechanism in general equilibrium.

Throughout the paper, $s_{i,t}$ can be interpreted as either the quantity of raw ideas or the

quality of an idea; these two interpretations are isomorphic.

Let $\tilde{J}_t$ denote the value of assigning one raw idea to a project:

$$\tilde{J}_t = \tilde{J}(q_t). \quad (3.2)$$

In this partial-equilibrium setting, $\tilde{J}_t$ is a function of the exogenous aggregate state q_t . Section 4 determines this value in general equilibrium.

Let $S_t(s_{i,t})$ denote the value of an innovator with raw-idea stock $s_{i,t}$. The innovator chooses between adjustment and nonadjustment:

$$S_t(s_{i,t}) = \max \{S_t^a(s_{i,t}), S_t^n(s_{i,t})\}, \quad (3.3)$$

where $S_t^a(s_{i,t})$ and $S_t^n(s_{i,t})$ denote the value if the innovator adjusts and does not adjust, respectively.

A launch assigns all accumulated raw ideas to the new project. The value of adjustment is

$$S_t^a(s_{i,t}) = -c + s_{i,t}\tilde{J}_t + E_t\Lambda_{t,t+1}S_{t+1}(\eta_{i,t+1}), \quad (3.4)$$

where $\Lambda_{t,t+1}$ is the stochastic discount factor. The first term is the setup cost, and the second is the value of the raw ideas assigned to the project. After launch, the raw-idea stock resets to zero and begins the next period with the new inflow $\eta_{i,t+1}$.

The value of nonadjustment is

$$S_t^n(s_{i,t}) = E_t\Lambda_{t,t+1}S_{t+1}(\phi s_{i,t} + \eta_{i,t+1}). \quad (3.5)$$

The solution has a threshold form.

Proposition 1 (Lumpy Innovation). *If $c > 0$, there exists a threshold $\bar{s}_t > 0$ such that innovators launch a project if and only if their stock of raw ideas exceeds $\bar{s}_t$.*


Proof. See Appendix A.1. □

Figure 2 visualizes Proposition 1. It plots the distribution of $s_{i,t}$ and the launch threshold $\bar{s}_t$. The shaded area is the mass of innovators that launch new projects.

Given the launch threshold, the law of motion for $s_{i,t+1}$ is

$$s_{i,t+1} = \phi s_{i,t} I\{s_{i,t} < \bar{s}_t\} + \eta_{i,t+1}, \quad (3.6)$$

Figure 2: Project Launches Along the Extensive Margin



where $I\{s_{i,t} < \bar{s}_t\}$ equals 1 if the innovator does not launch and 0 otherwise.

Aggregate project launches depend on the distribution of raw-idea stocks and the threshold $\bar{s}_t$. Aggregate shocks move the threshold, while past launches reshape the distribution. We next study the interaction between these two objects.

3.2 Fluctuations and the Distribution


Consider a favorable shock that raises the value of raw ideas, $\tilde{J}_t$. The higher value induces innovators to launch projects with smaller raw-idea stocks, shifting the launch threshold to the left (Figure 3a). More innovators then launch and reset their stocks. The subsequent distribution is therefore compressed, as shown by the red dashed line in Figure 3c. Lemma A.1(i) in the Online Appendix states this formally.

An adverse shock has the opposite immediate effect. A lower value of raw ideas shifts the launch threshold to the right (Figure 3b). Innovators delay project launch and carry their stocks forward, stretching the subsequent distribution. The green dashed line in Figure 3c illustrates this effect. Lemma A.1(ii) in the Online Appendix formalizes this stretching effect.


Threshold movements leave a persistent imprint on the distribution. A favorable shock induces launches and resets that cannot be undone when the threshold later returns. An adverse shock instead postpones launches and preserves accumulated stocks. The distribution therefore retains the history of past shocks, a property we call *distributional hysteresis*.

Distributional hysteresis is asymmetric when project launch is sufficiently lumpy. A


Figure 3: Aggregate Shocks and the Distribution of Raw-Idea Stocks




(a) A favorable shock shifts $\bar{s}_t$ left



(c) Asymmetric Effects on the Distribution



(b) An adverse shock shifts $\bar{s}_t$ right



(d) Moderation and the distribution of $s_{i,t}$

setup cost makes innovators wait until they have accumulated enough raw ideas to launch. When this cost is sufficiently high, reaching the launch threshold requires an inflow in the declining part of the idea-inflow distribution. For any given stock carried into the period, an inflow that leaves an innovator just below the threshold is therefore more likely than one that takes it just above. Across innovators, this produces a stationary density that declines at the launch threshold. As a result, a leftward movement in the threshold induces more launches and resets than an equal rightward movement postpones. Lemma 1 formalizes this asymmetry.

Lemma 1 (Asymmetric Distributional Hysteresis). *There exists a finite setup-cost cutoff $c^* > 0$ such that, if $c > c^*$, a favorable shock that lowers the launch threshold induces more additional project launches than an equally sized adverse shock postpones.*

Proof. See Appendix A.2. □

The condition $c > c^*$ is a sufficient condition to establish the asymmetry in Lemma 1. Why does it help? The density function of idea inflows can initially increase before declining. As a result, at a low launch threshold (small c), some innovators may be more likely to end up just above the threshold than just below it. This condition is not needed if one assumes a strictly decreasing density function.

The asymmetry extends to repeated threshold movements. For tractability, let the threshold alternate around $\bar{s}^*$:

$$\bar{s}_t = (1 + (-1)^t \sigma) \bar{s}^*.$$

Here $0 < \sigma < 1$ is the proportional threshold-fluctuation amplitude. In Section 4, stochastic aggregate shocks replace this deterministic sequence. Proposition 2 shows that repeated movements reduce the density at the launch threshold.

Proposition 2 (Fluctuations Reduce Density at the Launch Threshold). *There exists a finite setup-cost cutoff $c^* > 0$ such that, if $c > c^*$, then for every $\sigma \in (0, 1)$, the time-averaged stationary density f_σ of raw-idea stocks exists and satisfies*

$$f_\sigma(\bar{s}^*) < f^*(\bar{s}^*), \tag{3.7}$$

where f^* denotes the steady-state density in the absence of aggregate shocks.

Proof. See Appendix A.3. □

The cross-sectional distribution varies across the two phases of the threshold cycle. Averaging these distributions gives the time-averaged stationary distribution, which we refer to as the average distribution hereafter.

Figure 3d illustrates Proposition 2. Repeated threshold movements leave fewer innovators concentrated near the project-launch margin. Periods of lower aggregate volatility, which we call moderation, allow this concentration to build. The inherited distribution therefore changes how many project launches a subsequent adverse shock postpones.

Discussion Two features underlie asymmetric distributional hysteresis.

First, higher setup costs raise the launch threshold. Online Appendix C.2 examines how the degree of asymmetry varies with this threshold under alternative idea-inflow processes. At higher thresholds, the numerical examples show increasingly asymmetric launch responses to equal percentage movements in the threshold. This strengthens the asymmetry of distributional hysteresis highlighted above.

Second, project launch is one-sided. A favorable shock induces innovators to use their accumulated raw ideas, whereas an adverse shock postpones launches without reversing earlier ones. Online Appendix C.3 contrasts this structure with a capital-investment model in which firms can both invest and disinvest. A shift in the inaction region induces adjustment at one boundary while postponing it at the other. The distributional effects of these two responses can offset each other. In the appendix's calibrated example, aggregate fluctuations consequently have little effect on the average distribution of capital gaps.

4 The General Equilibrium Model

Section 3 isolates the project-launch mechanism in partial equilibrium. We now embed it in a medium-term business-cycle model that builds on Comin and Gertler (2006) and Comin et al. (2009). Aggregate conditions affect profits from new varieties and hence the value of assigning raw ideas to R&D projects, $\tilde{J}_t$. This value moves the launch threshold and changes the projects launched. Project building and adoption then translate these launches into new varieties, productivity, and output.

Innovators have heterogeneous raw-idea stocks and pay a fixed project-launch cost. They use skilled labor to generate raw ideas. Launched ideas then pass through project building and stochastic adoption before becoming intermediate-input varieties. Figure 4 summarizes this sequence.


The model has innovation, production, and business-cycle blocks. Section 4.1 completes the innovation block. Section 4.2 links new varieties to profits and productivity. Section 4.3 introduces aggregate shocks and monetary policy. Section 4.4 presents the market-clearing conditions and defines equilibrium. Section 4.5 discusses tractability and the model assumptions.

4.1 The Innovation Block

Phase I: Project Launch and Idea Creation Section 3 characterizes the project-launch decision in partial equilibrium. The general-equilibrium model makes two modifications.

First, the fixed setup cost is cA_t , where A_t is the number of existing intermediate-input varieties. Scaling the cost with A_t captures the rising complexity of launching an R&D project and prevents the launch friction from vanishing as the economy grows.

Figure 4: The Innovation Block



Second, idea inflows are endogenous.³ Innovator i hires skilled labor $N_{i,t}$ in period t to generate raw ideas in period $t + 1$:

$$\eta_{i,t+1} = \left(\zeta N_{i,t}^{\tilde{\zeta}} \right) \cdot \varepsilon_{i,t+1}^{\eta}, \quad (4.1)$$

where ζ governs idea productivity, ζ is the elasticity with respect to skilled labor, and $\varepsilon_{i,t+1}^{\eta}$ is a nonnegative idiosyncratic draw that is independent across innovators and time.

These modifications change the launch and waiting values in Equations (3.4) and (3.5). The value of launching is

$$S_t^a(s_{i,t}) = \max_{N_{i,t}} \left\{ -cA_t + s_{i,t} \tilde{J}_t + E_t \Lambda_{t,t+1} S_{t+1} \left(\zeta N_{i,t}^{\tilde{\zeta}} \cdot \varepsilon_{i,t+1}^{\eta} \right) - W_t^{RD} N_{i,t} \right\}, \quad (4.2)$$

where W_t^{RD} is the skilled-labor wage. The value of waiting is

$$S_t^n(s_{i,t}) = \max_{N_{i,t}} \left\{ E_t \Lambda_{t,t+1} S_{t+1} \left(\phi s_{i,t} + \zeta N_{i,t}^{\tilde{\zeta}} \cdot \varepsilon_{i,t+1}^{\eta} \right) - W_t^{RD} N_{i,t} \right\}. \quad (4.3)$$

The launch threshold and the law of motion for raw-idea stocks retain the form derived in Section 3.1.

³An earlier version of the paper treated idea inflows as exogenous. The qualitative results are unchanged under that specification, and the quantitative results are similar.

Phase II: Project Building and Adoption At launch, the innovator assigns all accumulated raw ideas to a pending R&D project. A fraction $\tilde{\lambda} < 1$ of pending projects is built each period. This time-to-build captures delays in securing funding, assembling research teams, and obtaining equipment (Bloom, 2007; Brown and Petersen, 2011; Chen et al., 2021).

Let J_t denote the value of an R&D idea in a built project. Pending ideas become built with probability $\tilde{\lambda}$; otherwise, they remain pending and survive at rate ϕ_A . Their value is therefore

$$\tilde{J}_t = \tilde{\lambda}J_t + \phi_A(1 - \tilde{\lambda})E_t\Lambda_{t,t+1}\tilde{J}_{t+1}. \quad (4.4)$$

Once a project is built, each R&D idea is adopted as a new variety with probability λ . Adoption requires $\bar{n}$ units of skilled labor, while ideas not adopted survive at rate ϕ_A . The value of a built R&D idea is

$$J_t = -W_t^{RD}\bar{n} + \lambda E_t\Lambda_{t,t+1}A_{t+1}V_{t+1} + (1 - \lambda)\phi_A E_t\Lambda_{t,t+1}J_{t+1}, \quad (4.5)$$

The three terms are the skilled-labor cost of adoption, the expected value of successful adoption, and the continuation value of an idea that remains unadopted.

Phase III: New Varieties and Profits An adopted idea becomes a new intermediate-input variety. Following Acemoglu et al. (2012), we assume that the innovator receives a one-period patent and sells its production right to an intermediate-input producer. The patent value equals the profit generated by the variety:

$$V_t = \Phi_t, \quad (4.6)$$

where Φ_t denotes the profit generated by one intermediate-input variety. Section 4.2 derives this profit from the intermediate-input producer's problem. The one-period patent is a tractability assumption; Online Appendix C.4 discusses its role and alternative specifications.

The value of a new variety and the skilled-labor cost of adoption jointly determine J_t and $\tilde{J}_t$, linking aggregate conditions back to the project-launch decision in Phase I.

The new varieties created in the innovation block contribute to the aggregate number of varieties A_t , whose evolution we describe next.

From Project Launches to Aggregate Technology Let $\tilde{x}_{i,t}$ denote R&D ideas in pending projects and $x_{i,t}$ those in built projects. When innovator i launches, all of its raw ideas enter a new pending project:

$$\tilde{x}_{i,t}^{new} = s_{i,t} I\{s_{i,t} \geq \bar{s}_t\}. \quad (4.7)$$

Let $\tilde{X}_t \equiv \int_i \tilde{x}_{i,t} di$ denote the aggregate stock of pending R&D ideas. It evolves according to

$$\tilde{X}_{t+1} = \phi_A(1 - \tilde{\lambda})\tilde{X}_t + \int_i \tilde{x}_{i,t+1}^{new} di, \quad (4.8)$$

The first term contains surviving ideas in projects that remain pending; the second aggregates ideas assigned to newly launched projects.

The aggregate stock of R&D ideas in built projects, X_t , evolves as

$$X_{t+1} = \phi_A(1 - \lambda)X_t + \tilde{\lambda}\tilde{X}_{t+1}. \quad (4.9)$$

Adoption converts built R&D ideas into new varieties. The aggregate number of varieties evolves according to

$$A_{t+1} = \phi_A A_t + \lambda X_t A_{t+1}. \quad (4.10)$$

where existing varieties survive at rate ϕ_A .

Equations (4.7)–(4.10) provide the link from the distribution of raw-idea stocks to aggregate technology. Aggregate conditions move $\bar{s}_t$, the distribution determines the raw ideas assigned to new projects, and project building and adoption determine how these launches change A_t . The production block, which we describe next, maps A_t into productivity and output.

4.2 The Production Block

The production block determines how new varieties affect output and the return to innovation. An expansion in A_t raises productivity by enlarging the set of intermediate inputs, while the profits generated by these inputs determine the value of a new variety.

Intermediate Goods Producer A competitive intermediate goods producer combines unskilled labor L_t with a composite intermediate input M_t . The composite aggregates A_t

differentiated inputs:

$$M_t = \left(\int_0^{A_t} M_t(j)^{\frac{1}{\vartheta}} dj \right)^{\vartheta}, \quad (4.11)$$

where $\vartheta > 1$ governs substitution across varieties. As in [Romer \(1990\)](#), an expansion in A_t raises productivity. Cost minimization yields demand for variety j :

$$M_{t,j} = \left(\frac{P_{t,j}^M}{P_t^M} \right)^{\frac{\vartheta}{1-\vartheta}} M_t, \quad (4.12)$$

where $P_t^M \equiv \left(\int_0^{A_t} (P_{t,j}^M)^{\frac{1}{1-\vartheta}} dj \right)^{1-\vartheta}$ is the price of the composite. The producer transforms L_t and M_t into intermediate goods according to

$$Y_{t,m} = L_t^{1-\gamma} M_t^{\gamma}, \quad (4.13)$$

and solves

$$\max_{M_t, L_t} P_{t,m} Y_{t,m} - W_t L_t - P_t^M M_t, \quad (4.14)$$

where $P_{t,m}$ is the price of intermediate goods and W_t is the unskilled-labor wage. The first-order conditions are

$$\gamma \left(\frac{L_t}{M_t} \right)^{1-\gamma} = \frac{P_t^M}{P_{t,m}}, \quad (1-\gamma) \left(\frac{M_t}{L_t} \right)^{\gamma} = \frac{W_t}{P_{t,m}}. \quad (4.15)$$

Intermediate Input Producers Each variety j is supplied by a monopolistically competitive producer that converts intermediate goods into $M_{t,j}$ one for one:

$$M_{t,j} = f(Y_{t,m}) = Y_{t,m}. \quad (4.16)$$

Facing demand (4.12), the producer charges a constant markup over marginal cost, $P_{t,j}^M = \vartheta P_{t,m}$. Its profit is

$$\Phi_{t,j} = (P_{t,j}^M - P_{t,m}) \left(\frac{P_{t,j}^M}{P_t^M} \right)^{\frac{\vartheta}{1-\vartheta}} M_t. \quad (4.17)$$

Symmetry implies $\Phi_t = \Phi_{t,j}$ for every j . Because $V_t = \Phi_t$, these profits feed back into adoption and project launch. Innovation in turn expands A_t , raising productivity and output through (4.11) and (4.13).

4.3 Business Cycle Block

The business cycle block introduces aggregate demand shocks, nominal price rigidity, and monetary policy. It also determines household demand and the supply of both types of labor.

The Household A representative household consumes, trades risk-free bonds Q_t in zero net supply, and supplies unskilled labor L_t and skilled labor N_t . Unskilled labor produces intermediate goods, while skilled labor is used for idea creation and adoption. Let W_t and W_t^{RD} denote the corresponding nominal wages, R_t the return on bonds, and Ξ_t aggregate dividends. The household solves

$$\max_{\{C_{t+i}, L_{t+i}, N_{t+i}, Q_{t+i}\}} E_t \sum_{i=0}^{\infty} \beta^i q_{t+i} \left[\frac{C_{t+i}^{1-\sigma}}{1-\sigma} - \frac{L_{t+i}^{1+\varphi}}{1+\varphi} - \left(\frac{N_{t+i}^{1+\eta}}{1+\eta} \right) \right] \quad (4.18)$$

subject to

$$P_t C_t = W_t L_t + W_t^{RD} N_t + Q_t - R_t^{-1} Q_{t+1} + \Xi_t \quad (4.19)$$

where the demand shock follows $\log q_{t+1} = \rho_q \log q_t + \sigma_q e_{t+1}^q$, with $e_{t+1}^q \sim N(0, 1)$. The Euler equation is

$$1 = E_t \left[\Lambda_{t,t+1} \Pi_{t+1}^{-1} R_t \right] \quad (4.20)$$

where $\Pi_{t+1} \equiv P_{t+1}/P_t$ and $\Lambda_{t,t+1} = \beta(C_t/C_{t+1})^\sigma (q_{t+1}/q_t)$. Labor supply satisfies

$$C_t^\sigma L_t^\varphi = \frac{W_t}{P_t}, \quad C_t^\sigma N_t^\eta = \frac{W_t^{RD}}{P_t}. \quad (4.21)$$

Retailers and Final Goods Nominal rigidity arises from quadratic price adjustment costs. A competitive final-good producer aggregates differentiated consumption goods

$y_{t,h}$:

$$Y_t = \left(\int_0^1 y_{t,h}^{\frac{1}{1-\mu}} dh \right)^{\mu} \quad (4.22)$$

where μ governs substitution across goods. The corresponding price index is

$$P_t = \left(\int p_{t,h}^{\frac{1}{1-\mu}} \right)^{1-\mu}.$$

Cost minimization implies

$$y_{t,h} = \left(\frac{p_{t,h}}{P_t} \right)^{\frac{\mu}{1-\mu}} Y_t. \quad (4.23)$$

A unit mass of monopolistically competitive retailers converts intermediate goods into differentiated consumption goods one for one:

$$y_{t,h} = Y_{t,m}. \quad (4.24)$$

Retailer h sets $p_{t,h}$ subject to the adjustment cost $\frac{\kappa}{2}(p_{t,h}/p_{t-1,h} - 1)^2 Y_t$. Its problem is

$$\max_{\{p_{t,h}\}} E_t \sum_j \Lambda_{t,t+j} \left[(p_{t+j,h} - P_{t,m}) \left(\frac{p_{t+j,h}}{P_{t+j}} \right)^{\frac{\mu}{1-\mu}} Y_{t+j} - \frac{\kappa}{2} \left(\frac{p_{t+j,h}}{p_{t+j-1,h}} - 1 \right)^2 Y_{t+j} \right],$$

subject to (4.23), where $\Lambda_{t,t+j} \equiv \prod_{k=0}^{j-1} \Lambda_{t+k,t+k+1}$. Defining $\pi_t \equiv P_t/P_{t-1} - 1$, the log-linearized price-setting condition gives the New Keynesian Phillips curve:

$$\pi_t = \beta E_t \pi_{t+1} + \frac{\mu}{\kappa(1-\mu)} \left(\frac{1}{\mu} - \frac{P_{t,m}}{P_t} \right). \quad (4.25)$$

Monetary Policy The central bank sets the nominal interest rate according to

$$R_t = \frac{1}{\beta} \frac{A_{t+1}}{A_t} \Pi_t^{\varphi_\pi}. \quad (4.26)$$

The term A_{t+1}/A_t tracks the efficient real interest rate. In the model, this rate moves with productivity growth (Wicksell, 1898). The coefficient φ_π governs the response to inflation. This specification also delivers the tractability discussed in Section 4.5.

4.4 Market Clearing and Equilibrium

The final- and intermediate-goods markets clear when

$$Y_t = C_t, \quad Y_{t,m} = \int_0^1 y_{t,h} dh + \int_0^{A_t} M_{t,j} dj. \quad (4.27)$$

The unskilled- and skilled-labor markets clear when labor demand from production and innovation equals household supply. Risk-free bonds are in zero net supply, so $Q_t = 0$.

Agent optimality, the laws of motion for the innovation distributions and aggregate technology, the monetary-policy rule, and these market-clearing conditions jointly define a recursive competitive equilibrium. Online Appendix C.5 gives the formal definition and describes the solution method.

4.5 Tractability and Model Assumptions

The model retains heterogeneous project-launch decisions and remains tractable. Without further restrictions, the distributions of raw ideas and R&D ideas would affect expected technology growth and skilled-labor demand. Both distributions would then become aggregate state variables.⁴

Two assumptions limit this feedback. First, the Taylor rule offsets expected productivity growth in the household Euler equation. Second, skilled-labor supply is perfectly elastic, so the distribution of projects affects skilled employment but not the skilled-labor wage. Aggregate conditions still affect innovation through patent profits, adoption costs, and discounting, while innovation continues to determine technology growth. The innovation distributions need not enter household consumption and labor-supply decisions. This structure allows us to solve the business-cycle and innovation blocks sequentially.

The model also abstracts from knowledge spillovers across innovators: a project launch does not directly expand other innovators' raw-idea stocks. This restriction isolates the selection-and-reset mechanism from channels through which current research expands future research opportunities (Arora et al., 2021).

Online Appendix C.4 discusses these restrictions. It also covers the assumptions governing patent duration, idea obsolescence, passive adoption, and time-to-build.

⁴These distributions interact with the kinked project-launch policy. Standard approximate-aggregation methods such as Krusell and Smith (1998) can therefore generate large errors in this setting.

5 Model Calibration and Estimation

We discipline the quantitative model in three stages. Section 5.1 combines parameters taken from the literature with a joint calibration of the innovation and aggregate blocks. Section 5.2 evaluates untargeted cyclical and distributional implications. Section 5.3 then recovers the historical demand shocks and model states used in the quantitative exercises.

5.1 Parameterization and Calibration

We calibrate the model to the U.S. economy at an annual frequency. The underlying firm-level data are the 1968–2020 Compustat sample described in Section 2. Aggregate time-series moments and the historical shock recovery use 1984–2018. We fix parameters that have standard counterparts in the literature and jointly calibrate the six parameters that govern long-run innovation, lumpy project initiation, and aggregate propagation.

Fixed Parameters The discount factor β implies an annual interest rate of 4%, and $\varphi = 0.5$ gives an inverse Frisch elasticity for unskilled labor of 2. We set $\sigma = 1$ and $\eta = 0$ as in the tractable specification described in Section 4. In the New Keynesian block, $\kappa = 90$ and $\mu = 1.1$ imply a Phillips-curve slope of 0.1, following Ottonello and Winberry (2020). We set the Taylor-rule coefficient to $\varphi_\pi = 1.69$ based on our estimation and the annual persistence of the demand shock to $\rho_\ell = 0.4$, corresponding to a quarterly persistence of 0.8.

For innovation and production, we set the intermediate-input share to $\gamma = 0.5$ following Comin and Gertler (2006). We set $\vartheta = 2$, so that $\vartheta\gamma = 1$, following Kung and Schmid (2015). The survival rate of intermediate inputs is $\phi_A = 0.97$, implying annual depreciation of 3%. The adoption probability is $\lambda = 0.2$, corresponding to an average adoption time of five years (Anzoategui et al., 2019). Finally, we set the elasticity of idea creation with respect to R&D expenditure to $\zeta = 0.5$ following Acemoglu et al. (2018).

Calibration Strategy We jointly calibrate ζ , c , $1 - \phi$, $\tilde{\lambda}$, σ_ℓ , and $\bar{n}$ to the moments in Table 2. These moments discipline three parts of the model.

First, average GDP growth disciplines long-run innovation. Idea inflow is $\eta_{i,t+1} = \zeta N_{i,t}^\zeta \varepsilon_{i,t+1}^\eta$, where $\varepsilon_{i,t+1}^\eta$ follows a $\mathcal{N}(1, 1)$ distribution truncated below at zero. A higher ζ raises idea creation and, through project launch and adoption, the average growth rate of technology and output.

Second, firm-level R&D moments discipline project initiation and completion. The average duration between spikes informs the fixed setup cost c : a higher setup cost raises the launch threshold and lengthens the time between large adjustments. The max-to-mean ratio of R&D expenditure informs the obsolescence rate $1 - \phi$. Faster obsolescence makes large accumulations of unused raw ideas less likely and compresses the size of R&D spikes. The autocorrelation of the spike rate informs the project-completion probability $\tilde{\lambda}$, which governs how long launched projects remain under development.

Finally, the standard deviation and autocorrelation of GDP growth jointly discipline σ_ϵ and $\bar{n}$. The former determines the volatility of demand shocks. The latter scales the skilled labor required by an R&D project and thereby affects how aggregate conditions translate into project values and subsequent technology growth. Demand shocks generate relatively transitory output movements, whereas innovation produces more persistent TFP movements. The volatility and persistence of GDP growth therefore help separate the two parameters.

Calibrated Parameters Table 2 reports the calibrated parameters and moments. The annual obsolescence rate of unused raw ideas is 15%. The estimate $\tilde{\lambda} = 0.40$ implies that an R&D project takes 2.5 years on average to move from launch to completion. The estimate $c = 0.001$ implies a setup cost equal to approximately 8% of the value of the raw ideas assigned to a project. The model also closely matches the targeted moments in the table.

5.2 Untargeted Model Implications

The calibration targets the timing and size of R&D spikes and the aggregate dynamics of GDP. We evaluate the model using two implications that are not targeted: the cyclicity of large R&D adjustments and the cross-sectional distribution of adjustment timing.

Cyclicity of Large R&D Adjustments Section 2 documents that the R&D spike rate is procyclical. Neither its correlation with GDP growth nor the cyclicity of aggregate R&D is targeted in the calibration. The last two rows of Table 2 show that the model reproduces both moments closely. The correlation of the spike rate with GDP growth is 0.52 in both the model and the data, while the correlation of aggregate R&D with log GDP is 0.68 in the model and 0.62 in the data. Thus, aggregate conditions move the project-launch margin with realistic cyclicity.

Distribution of Innovators The model's key quantitative state is the distribution of raw-idea stocks, $s_{i,t}$, which is not directly observable. We therefore evaluate this distribution

Table 2: Calibration: Parameters and Moments

<i>Panel A: Calibrated Parameters</i>		
Parameter	Description	Value
ζ	Innovation productivity	0.46
c	Fixed R&D setup cost	$1e-3$
$1 - \phi$	Raw-idea obsolescence	0.15
$\tilde{\lambda}$	Project completion probability	0.40
$\bar{n}$	Skilled labor per project	0.79
σ_ϱ	Standard deviation of demand shocks	0.02
<i>Panel B: Data and Model Moments</i>		
Moment	Data	Model
<i>Targeted moments</i>		
Mean GDP growth	0.028	0.028
Mean duration between R&D spikes	6.89	7.08
Max-to-mean ratio of R&D expenditure	2.40	2.41
AR(1) of the R&D spike rate	0.57	0.51
AR(1) of GDP growth	0.36	0.37
Standard deviation of GDP growth	0.017	0.018
<i>Untargeted moments</i>		
Correlation of the R&D spike rate with GDP growth	0.52	0.52
Correlation of aggregate R&D with log GDP	0.62	0.68


Notes: The firm-level moments use the 1968–2020 Compustat sample described in Section 2. Aggregate time-series moments use 1984–2018. Spike-based moments pool definitions based on R&D increases from 40% to 100%. The max-to-mean ratio is the average firm-level ratio of maximum to mean R&D expenditure. Mean duration between spikes is the average number of years between two spike events for each firm.

using the time elapsed since an innovator’s last R&D spike, denoted by $TE_{i,t}$. A project launch resets the raw-idea stock, after which the innovator begins accumulating ideas again. The distribution of elapsed time therefore contains information about where innovators lie in this accumulation cycle. Lemma C.1 in Online Appendix C.6 formalizes the relationship between this observable and the distribution of raw-idea stocks.

Figure 5a compares the empirical distribution of $TE_{i,t}$ with its model counterpart. Although this distribution is not targeted, the model closely matches its overall shape. Through Lemma C.1, this agreement provides discipline on the latent distribution of raw-idea stocks used in the quantitative analysis.

The adjustment hazard helps assess whether time elapsed since the last R&D spike is informative about proximity to project launch. In the model, innovators that have waited longer since their previous launch have had more time to accumulate raw ideas, making

Figure 5: Untargeted Distributional Implications



Notes: Empirically, *time elapsed since the last spike* is measured as the interval between the current spike episode and the most recent one. Due to amortization and the time-to-build nature of R&D projects (averaging 2.5 years in the model), firms' R&D expenses may exhibit consecutive spikes, of which we retain only the first to avoid misidentifying multiple spikes from the same project. Adjustment hazard is calculated as the share of firms that spike, conditional on surviving to a given number of years since the previous spike.

another launch more likely. We compute the empirical adjustment hazard as the probability of an R&D spike at each elapsed duration, conditional on no intervening spike. Figure 5b shows that the empirical and model-implied hazards are close in both slope and level. This agreement supports the model's prediction and the use of elapsed time as an indicator of proximity to project launch in the firm-level evidence in Section 7.2.

5.3 Estimating Aggregate Shocks


With the calibrated parameters fixed, we estimate the sequence of demand shocks, ε_t^q , by fitting the model to annual U.S. GDP growth from 1984 to 2018. These shocks, together with the distributions they generate, provide the initial conditions for the quantitative exercises in Section 6.⁵

Figure A.1a compares model-implied and observed GDP growth. The model tracks the observed series closely. The figure also reports 80% and 99% one-period-ahead prediction bands, constructed from the model state in the preceding year.

Figure A.1b provides an external validation that is not used in estimating the shocks. It compares the model-implied growth of technology, A_t , with the utilization-adjusted

⁵We begin the simulation in 1979 and treat the first five years as a burn-in period. The estimation minimizes the mean squared deviation between model-implied and observed GDP growth, augmented by a penalty for unlikely shock sequences.

Figure 6: Great Recession: Distribution Matters



Notes: Panel (a) compares the estimated distribution of the raw-idea stock $s_{i,t}$ immediately before the 2009 shock with the average cross-sectional distribution implied by the model. Panel (b) displays the cumulative response of TFP growth to the estimated 2008 and 2009 shocks. The red curve uses the historical pre-recession distribution, while the blue curve uses the average distribution. Both exercises begin from the same aggregate demand state and impose the same recessionary shocks.

TFP growth series of [Fernald \(2014\)](#). The model captures the medium-run movement in measured TFP growth, although its implied series is smoother.⁶ [Online Appendix A](#) discusses the model's interpretation of the evolution of TFP.

6 Distributional Hysteresis: Propagation, Policy, and Growth

With the calibrated parameters and recovered historical shocks in hand, we now quantify how the inherited distribution of raw ideas shapes aggregate dynamics. [Section 6.1](#) shows how the pre-recession distribution amplified the effects of the Great Recession. [Section 6.2](#) traces this vulnerable distribution to the preceding period of moderation and then studies the resulting trade-off for stabilization policy. [Section 6.3](#) examines how the same distributional mechanism affects average growth.

6.1 Great Recession: Distribution Matters

How much did the inherited distribution of raw ideas amplify the Great Recession? We answer this question by applying the estimated recessionary shocks to two economies:

⁶In the model, TFP moves only through technological progress. Measured TFP also reflects factors such as misallocation ([Hsieh and Klenow, 2009](#)), changes in workforce composition ([Moscarini and Postel-Vinay, 2018](#)), and measurement error.

one with the historical pre-recession distribution and another with the average distribution implied by the model. Because the model is estimated using annual U.S. GDP growth, the largest contractionary shock occurs in 2009. We therefore recover the distribution of $s_{i,t}$ immediately before that shock.

The pre-recession distribution was unusually vulnerable to an adverse shock. Figure 6a compares the estimated 2009 distribution of raw-idea stocks with the average distribution. The pre-recession distribution places substantially more mass near the project-launch threshold. When the contraction raises this threshold, more innovators postpone project launch and leave their accumulated ideas outside active R&D projects.

Figure 6b isolates this distributional amplification by imposing the same 2008 and 2009 shocks on both economies. The effects unfold gradually as project building and adoption transmit the initial launch shortfall to technology. Under the historical distribution, the cumulative loss in TFP growth reaches approximately 4.7% in the long run. Under the average distribution, the same shocks produce a loss of about 3.3%. The inherited distribution therefore amplifies the long-run TFP loss from the Great Recession shocks by approximately 40% relative to the average-distribution counterfactual.


6.2 The Great Twins: Great Moderation and Great Recession

What produced the vulnerable distribution inherited in 2009? We argue that the moderation of aggregate fluctuations during 2002–2007 allowed more innovators to accumulate raw ideas near the project-launch threshold. This inherited distribution then amplified the effects of the Great Recession shocks.

To quantify this channel, we construct counterfactual histories for 2002–2007 by adding mean-zero disturbances to the estimated demand shocks from 2003 to 2007. Each history generates a different realized pre-recession volatility and, through project launches and resets, a different distribution of raw-idea stocks. Starting from each resulting distribution, we then impose the same estimated 2008 and 2009 shocks. This exercise isolates how the preceding history changes the propagation of a common recessionary shock sequence.

Figure 7 shows that increasing realized volatility during 2002–2007 by 20% reduces the 40-year cumulative TFP loss caused by the Great Recession shocks by 1.5 percentage points. Greater pre-recession volatility induces more project launches and resets, leaving fewer innovators near the launch threshold when the recession begins. Conversely, the moderation following the dot-com recession allowed this mass to build and made the economy more vulnerable to the subsequent contraction. This is the Great Twins mecha-

Figure 7: The Great Twins



Notes: The figure reports the 40-year cumulative response of TFP growth to the estimated 2008 and 2009 shocks across counterfactual economies with different realized pre-recession volatilities. We generate these histories by setting $\tilde{\varepsilon}_t^q = \varepsilon_t^q + \tilde{\sigma}\tilde{\varepsilon}_t$ from 2003 to 2007, where $\tilde{\varepsilon}_t$ is drawn from a standard normal distribution. The horizontal axis reports the average percentage increase in the realized volatility of q_t over 2003–2007. The vertical axis reports the average 40-year cumulative response of TFP growth to the same 2008 and 2009 shocks.


nism.

This experiment corresponds to the “Good Luck” interpretation of the Great Moderation emphasized by [Justiniano and Primiceri \(2008\)](#). The calibrated shock process remains unchanged; the counterfactual histories differ only in their realized shocks before 2008. The result therefore captures how lower realized volatility changes the inherited distribution even when the economy’s underlying exposure to aggregate shocks is unchanged.

Policy Implication: The Tail-Risk–Variance Trade-off The Great Moderation may also have reflected improved stabilization policy rather than good luck. Under this “Good Policy” view, monetary policy reduced aggregate volatility through a stronger response to inflation, represented here by a larger φ_π in the Taylor rule ([Galí and Gambetti, 2009](#)). In our model, this policy-induced reduction in volatility leaves more innovators close to launching projects. Subsequent contractions that monetary policy cannot offset therefore postpone more launches and cause larger losses.

We evaluate this implication in two stages. First, we simulate stationary economies under alternative values of φ_π . A larger φ_π lowers unconditional GDP volatility and generates a different stationary distribution of raw-idea stocks. Second, we expose each economy to the same one-standard-deviation adverse aggregate demand shock and set $\varphi_\pi = 0$ throughout the recessionary period, mimicking a contraction that monetary policy cannot offset at the zero lower bound. Holding the policy response fixed during the

Figure 8: Policy Implications of the Great Twins



Notes: Panel (a) reports the 40-period cumulative loss in TFP growth following a one-standard-deviation adverse aggregate demand shock across economies previously simulated under different values of φ_π . During the recessionary period, φ_π is set to zero, so monetary policy cannot offset the shock. The upper horizontal axis reports the pre-recession policy parameter, and the lower horizontal axis reports the associated standard deviation of GDP growth. Panel (b) plots the stationary distribution of the raw-idea stock $s_{i,t}$ generated by each pre-recession policy regime.

contraction isolates the effects of stabilization policy before the shock, which operates through the inherited distribution of raw-idea stocks. Thus, φ_π on the upper horizontal axis of Figure 8a indexes the pre-recession policy regime.


Figure 8a reports the resulting tail-risk-variance trade-off. A larger pre-recession φ_π lowers the standard deviation of GDP growth but increases the cumulative TFP loss caused by the common adverse shock. Stronger stabilization therefore lowers unconditional GDP volatility but increases the losses from a severe contraction that monetary policy cannot offset.

The inherited distribution generates this trade-off. Figure 8b shows that a larger φ_π shifts the stationary distribution toward larger raw-idea stocks. When the adverse shock raises the launch threshold, more innovators postpone projects and leave accumulated ideas outside active R&D. Stabilization policy therefore affects the propagation of future shocks through its effect on the distribution of innovators.

6.3 Great Moderation and Growth

Stabilization policy also affects average growth. A larger φ_π reduces aggregate volatility but leaves more raw ideas outside active R&D projects, as shown in Figure 8b. Because

Figure 9: Great Moderation and Growth



Notes: The blue dashed line plots average output growth across different values of the monetary policy parameter φ_π . The gray dot-dash line plots the same relationship when idea inflows are fixed at their average across all states in the baseline model, shutting down the cyclical idea-creation channel while preserving the distributional utilization channel.

unused raw ideas depreciate, delaying project launch causes more ideas to lose value before they enter active R&D. The blue line in Figure 9 reports the resulting trade-off: a larger φ_π (upper horizontal axis) produces lower output volatility (lower horizontal axis) and lower average output growth (vertical axis).

Fluctuations affect growth through both the utilization of accumulated ideas and the creation of new ideas. To isolate the first channel, we fix idea inflows at their average level across aggregate states. The gray dot-dash line in Figure 9 shows that the negative relationship between the strength of stabilization and growth remains. The growth effect therefore operates even when aggregate conditions no longer affect the inflow of new ideas.

The literature offers competing predictions about the relationship between volatility and long-run growth. Schumpeterian mechanisms emphasize how recessions can accelerate creative destruction, whereas uncertainty and financial frictions can cause volatility to discourage innovation (Jones et al., 2000; Aghion et al., 2010). Our mechanism operates through a different channel. Moderation delays the utilization of accumulated ideas, and obsolescence converts this delay into lower technological progress. Empirical estimates are mixed (Kormendi and Meguire, 1985; Ramey and Ramey, 1995; Judson and Orphanides, 1999), but tend to be negative in developing economies and positive in advanced economies (Bakas et al., 2019). The positive relationship estimated for advanced economies is consistent with our model's prediction.

When Ideas Are Getting Harder to Find Recent evidence suggests that ideas have become harder to find over the past 40 years (Bloom et al., 2020; Kogan et al., 2017). Online Appendix C.7 shows that lowering the average inflow of ideas increases the real GDP loss caused by the Great Recession. The distributional mechanism therefore becomes more consequential as the production of new ideas becomes more difficult. If this trend continues, the mechanism we emphasize will become increasingly relevant in the future.

7 Empirical Evidence

We now examine the model’s aggregate prediction and underlying mechanism. Section 7.1 uses cross-country data to examine whether greater moderation is associated with deeper subsequent recessions. Section 7.2 uses firm-level data on time elapsed since the last R&D spike to assess the predicted pre-recession concentration of innovators near the project-launch threshold. Section 7.3 shows that R&D spikes predict subsequent patenting, linking large R&D increases to innovation outcomes.

7.1 The Great Twins in the Data

Output losses after the 2008 recession varied sharply across countries. Ball (2014) estimates losses ranging from negligible levels in Australia and Switzerland to more than 30% in Greece, Hungary, and Ireland. Is this variation related to the degree of moderation before the recession, as implied by our model? We examine this Great Twins relationship using cross-country data.


For each country, we measure moderation as the decline in GDP-growth volatility during 1995–2007 relative to its longer pre-2008 history. Recession depth is the gap between actual GDP and the level implied by the pre-recession growth trend.⁷

Panel (a) of Figure 10 presents the baseline relationship. OECD economies that experienced greater moderation before 2008 subsequently suffered deeper recessions. The correlation between moderation and recession depth is 0.64 ($p = 0.00$).

The relationship documented so far is a correlation. It remains positive and statistically significant after controlling for cross-country differences in real interest rates and pre-recession credit expansion (Online Appendix B.3, Table B.2). To obtain evidence on the mechanism, we use its prediction that the relationship should be stronger in economies

⁷Online Appendix B.3 provides the formal definition of the GDP gap. The note to Figure 10 describes the normalization used in the broader-sample comparisons.

Figure 10: The Great Twins across Countries



Notes: Each dot represents a country. GDP data are from the World Bank. Moderation and recession depth are defined in the text. Panel (a) uses the OECD sample; the correlation is 0.64 ($p = 0.00$). Panels (b) and (c) use the broader World Bank sample, with recession depth normalized by pre-recession trend growth, measured by the average growth rate over the preceding fifteen years. Panel (b) contains 26 countries with an R&D-to-GDP ratio above 1%; the correlation is 0.61 ($p = 0.00$). Panel (c) contains 39 countries with a ratio below 1%; the correlation is -0.06 ($p = 0.70$).

where innovation contributes more to aggregate growth. Panels (b) and (c) of Figure 10 therefore divide the broader World Bank sample by R&D intensity. Among countries with an R&D-to-GDP ratio above 1%, greater pre-recession moderation is associated with a deeper subsequent recession: the correlation is 0.61 ($p = 0.00$). Among countries with an R&D-to-GDP ratio below 1%, the relationship disappears: the correlation is -0.06 ($p = 0.70$). This heterogeneity is consistent with the innovation mechanism underlying the Great Twins.

Additional Results Online Appendix [B.3](#) considers alternative R&D-intensity cutoffs, measures recession depth using TFP rather than real GDP, and extends the analysis beyond 2008 to recessions in 18 advanced economies since World War II. The positive relationship between moderation and recession depth remains across these exercises.

7.2 Firm-Level Evidence on the Distributional Mechanism

The Great Twins mechanism operates through the pre-recession distribution of innovators around the project-launch threshold. Its key empirical implication is that moderation before 2008 left an unusually large mass of innovators close to launching. Raw-idea stocks are not directly observable. As discussed in Section [5.2](#), the empirical adjustment hazard increases with time elapsed since the last R&D spike (Figure [5b](#)). The model closely matches this hazard in both slope and level. Together, these findings support using the distribution of elapsed time as a proxy for the distribution of raw-idea stocks. We therefore use this observable distribution to examine whether the onset of the Great Recession is marked by the concentration predicted by the mechanism.


Figure [11](#) compares the 2008 distribution with the corresponding average distribution in the data and the model. In both panels, the 2008 distribution places greater mass around five to eight years since the previous spike. This range contains the average 6.05-year interval between spikes. The agreement between the firm-level data and the distribution recovered from GDP growth provides evidence for the model’s distributional mechanism: at the onset of the recession, more innovators were approaching the usual timing of their next project launch.

7.3 R&D Spikes Predict Subsequent Patenting

We finally examine whether R&D spikes predict subsequent innovation output. In the model, project launches lead to new technologies through project building and adoption. If R&D spikes capture these launches, they should predict subsequent patenting. This evidence complements the procyclicality of large R&D adjustments documented in Section [2](#).

Data Used in This Subsection To connect R&D spending to innovation outcomes, we link Compustat firms to USPTO patents following [Arora et al. \(2021\)](#). We measure innovation quantity by the number of new patents and consider three measures of patent impact: citations, the stock-market-based patent value of [Kogan et al. \(2017\)](#), and the text-

Figure 11: Time Elapsed since the Last R&D Spike in 2008



Notes: Panel (a) compares the empirical distribution in 2008 with its average over the sample. Panel (b) compares the model-implied distribution in 2008, recovered by fitting the model to U.S. GDP growth, with the average distribution in model simulations. Elapsed time is truncated at 20 years in both panels.

based importance measure of [Kelly et al. \(2021\)](#), which is higher for patents that are both novel and influential.

Econometric Method We examine whether an R&D spike predicts subsequent patenting beyond the linear relationship with the size of the R&D change. For each horizon h , we estimate


$$\Delta_h y_{i,t+h} = \alpha_{i,h} + \beta_h \cdot \text{Spike}_{i,t} + \gamma_h \Delta \text{R\&D}_{i,t} + \theta_h X_{i,t} + \mu_t + u_{i,t,h}, \quad (7.1)$$

Here, $y_{i,t+h}$ denotes one of the four patent outcomes. To reduce the sparsity of annual patenting, we smooth each outcome using a three-year backward moving average. The dependent variable measures the change from t to $t + h$, divided by the average of its initial and final values.

The variable $\Delta \text{R\&D}_{i,t}$ measures the annual percentage change in R&D expenditure, and $\text{Spike}_{i,t}$ equals one when this change exceeds 40%. The terms $\alpha_{i,h}$ and μ_t are firm and year fixed effects. The controls $X_{i,t}$ include sales, total assets, acquisitions, and lagged R&D expenditure.

The two coefficients capture different features of R&D adjustment. The coefficient γ_h captures the linear association between the size of the R&D change and subsequent patenting. The coefficient β_h captures the additional association with an R&D spike beyond this linear relationship. A positive β_h therefore indicates that the spike indicator contains information about subsequent innovation outcomes beyond that captured by

Figure 12: R&D Spikes and Subsequent Patenting



Notes: Each panel reports cumulative estimates from the panel regression described in the text. The dark curve plots the coefficient on an R&D spike, defined as an increase greater than 40%, conditional on the continuous change in R&D expenditure and the controls described in the text. The orange curve plots the coefficient on the continuous R&D change, scaled to a 100% increase. To connect the estimates to the model, the legend labels these coefficients the extensive and intensive margins, respectively. Shaded areas denote 68% and 95% confidence intervals based on Driscoll-Kraay standard errors.

the continuous R&D-change term.

Results Figure 12 reports the estimates across horizons. The dark curves show the spike coefficients, while the orange curves show the continuous R&D-change coefficients scaled to a 100% increase. An R&D spike predicts sustained increases in all four patent outcomes. For patent counts and total citations, the spike coefficient is roughly twice the estimate associated with the continuous R&D increase. This ratio is larger for total patent value and technological importance. Thus, R&D spikes predict both more patents and greater total economic and technological impact beyond the linear relationship with expenditure changes.

Together, these exercises connect the model's aggregate Great Twins prediction to firm-level evidence on the timing and outcomes of lumpy R&D adjustment.

8 Conclusion

We develop a business-cycle model in which heterogeneous innovators accumulate raw ideas and launch projects at an endogenous threshold. Launches reset raw-idea stocks, making their cross-sectional distribution path dependent. When project launch is sufficiently lumpy, more innovators lie just below the launch threshold than just above it. Favorable threshold movements therefore induce more launches than equal adverse movements postpone. High aggregate volatility repeatedly moves the threshold and, through launches and resets, leaves fewer innovators near it, whereas moderation allows this mass to build. The inherited distribution then determines how strongly a contraction suppresses project launches, technology, and output.

This mechanism generates the Great Twins hypothesis. Holding fixed the shocks of the Great Recession, the estimated pre-recession distribution produces a long-run cumulative TFP loss of 4.7%, compared with 3.3% under the model's average distribution. The model traces this amplification to the moderation preceding 2008. It also reveals a policy trade-off: stronger stabilization lowers unconditional GDP volatility but increases the losses from a severe contraction that monetary policy cannot offset. Stronger stabilization also delays the use of accumulated ideas, lowering average growth when unused ideas become obsolete.

The evidence supports both the Great Twins relationship and its distributional mechanism. Firm-level R&D adjustments are episodic, the frequency of large increases is procyclical, and these increases predict subsequent patent counts and total patent impact. Across countries, greater moderation before 2008 is associated with deeper subsequent recessions. This correlation survives controls for real interest rates and pre-recession credit expansion and is concentrated among R&D-intensive economies, consistent with the model's innovation mechanism. At the firm level, the probability of an R&D spike increases with time since the previous spike, supporting the use of elapsed time as a proxy for proximity to project launch. The 2008 distribution of elapsed time closely matches its model counterpart, providing micro-level support for the mechanism operating through the pre-recession concentration of innovators near the launch threshold.

The broader lesson is that business-cycle stabilization changes not only current volatility but also the distribution governing future shock propagation. Accounting for this inherited distribution is therefore important for understanding medium-term business cycles and evaluating the effects of stabilization policy on vulnerability and growth.

References

- Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous**, “The environment and directed technical change,” *American economic review*, 2012, 102 (1), 131–166.
- , **Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr**, “Innovation, reallocation, and growth,” *American Economic Review*, 2018, 108 (11), 3450–91.
- Aghion, Philippe, George-Marios Angeletos, Abhijit Banerjee, and Kalina Manova**, “Volatility and growth: Credit constraints and the composition of investment,” *Journal of Monetary Economics*, 2010, 57 (3), 246–265.
- 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.
- Arora, Ashish, Sharon Belenzon, and Lia Sheer**, “Knowledge spillovers and corporate investment in scientific research,” *American Economic Review*, 2021, 111 (3), 871–898.
- Bachmann, Rüdiger, Ricardo J Caballero, and Eduardo M R A Engel**, “Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model,” *American Economic Journal: Macroeconomics*, 2013, 5 (4), 29–67.
- Bagliano, Fabio C and Claudio Morana**, “It ain’t over till it’s over: A global perspective on the Great Moderation-Great Recession interconnection,” *Applied Economics*, 2017, 49 (49), 4946–4969.
- Bakas, Dimitrios, Georgios Chortareas, and Georgios Magkonis**, “Volatility and growth: a not so straightforward relationship,” *Oxford Economic Papers*, 2019, 71 (4), 874–907.
- Baley, Isaac and Andrés Blanco**, “Aggregate dynamics in lumpy economies,” *Econometrica*, 2021, 89 (3), 1235–1264.
- Ball, Laurence**, “Long-term damage from the Great Recession in OECD countries,” *European Journal of Economics and Economic Policies: Intervention*, 2014, 11 (2), 149–160.
- Barlevy, Gadi**, “On the cyclicalities of research and development,” *American Economic Review*, 2007, 97 (4), 1131–1164.
- Benedetti-Fasil, Cristiana, Giammario Impullitti, Omar Licandro, Petr Sedláček, and Adam Hal Spencer**, “Heterogeneous Firms, Growth and the Long Shadows of Business Cycles,” 2025. Working paper.
- Benigno, Gianluca and Luca Fornaro**, “Stagnation traps,” *The Review of Economic Studies*,

2018, 85 (3), 1425–1470.

Berlingieri, Giuseppe, Maarten De Ridder, Danial Lashkari, and Davide Rigo, “Creative Destruction through Innovation Bursts,” Discussion Paper 2095, Centre for Economic Performance, London School of Economics 2026. Revised May 2026.

Bianchi, Francesco, Howard Kung, and Gonzalo Morales, “Growth, slowdowns, and recoveries,” *Journal of Monetary Economics*, 2019, 101, 47–63.

Bilbiie, Florin O, Fabio Ghironi, and Marc J Melitz, “Endogenous entry, product variety, and business cycles,” *Journal of Political Economy*, 2012, 120 (2), 304–345.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, “Are Ideas Getting Harder to Find?,” *American Economic Review*, 2020, 110 (4), 1104–1144.

Bloom, Nick, “Uncertainty and the Dynamics of R&D,” *American Economic Review*, 2007, 97 (2), 250–255.

Broda, Christian and David E Weinstein, “Product creation and destruction: Evidence and price implications,” *American Economic Review*, 2010, 100 (3), 691–723.

Brown, James R and Bruce C Petersen, “Cash holdings and R&D smoothing,” *Journal of Corporate Finance*, 2011, 17 (3), 694–709.

Caballero, Ricardo J and Eduardo MRA Engel, “Explaining investment dynamics in US manufacturing: a generalized (S, s) approach,” *Econometrica*, 1999, 67 (4), 783–826.

—, —, and **John C Haltiwanger**, “Plant-level adjustment and aggregate investment dynamics,” *Brookings papers on economic activity*, 1995, pp. 1–54.

Cerra, Valerie, Antonio Fatás, and Sweta C. Saxena, “Hysteresis and Business Cycles,” *Journal of Economic Literature*, 2023, 61 (1), 181–225.

Cette, Gilbert, John Fernald, and Benoit Mojon, “The pre-Great Recession slowdown in productivity,” *European Economic Review*, 2016, 88, 3–20.

Chen, Zhao, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Yi Xu, “Notching R&D investment with corporate income tax cuts in China,” *American Economic Review*, 2021, 111 (7), 2065–2100.

Comin, Diego A, Mark Gertler, and Ana Maria Santacreu, “Technology innovation and diffusion as sources of output and asset price fluctuations,” Technical Report, National Bureau of Economic Research 2009.

Comin, Diego and Mark Gertler, “Medium-term business cycles,” *American Economic Review*, 2006, 96 (3), 523–551.

Cooper, Russell W and John C Haltiwanger, “On the nature of capital adjustment costs,” *The Review of Economic Studies*, 2006, 73 (3), 611–633.

- Fatas, Antonio**, "Do business cycles cast long shadows? Short-run persistence and economic growth," *Journal of economic growth*, 2000, 5, 147–162.
- Fernald, John**, "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Technical Report Working Paper 2014-04, Federal Reserve Bank of San Francisco 2014.
- Galí, Jordi and Luca Gambetti**, "On the sources of the great moderation," *American Economic Journal: Macroeconomics*, 2009, 1 (1), 26–57.
- Guerron-Quintana, Pablo A and Ryo Jinnai**, "Financial frictions, trends, and the great recession," *Quantitative Economics*, 2019, 10 (2), 735–773.
- Hsieh, Chang-Tai and Peter J Klenow**, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, 124 (4), 1403–1448.
- Jones, Larry E, Rodolfo Manuelli, and Ennio Stacchetti**, "Technology (and policy) shocks in models of endogenous growth," 2000.
- Judson, Ruth and Athanasios Orphanides**, "Inflation, volatility and growth," *international Finance*, 1999, 2 (1), 117–138.
- Justiniano, Alejandro and Giorgio E Primiceri**, "The time-varying volatility of macroeconomic fluctuations," *American Economic Review*, 2008, 98 (3), 604–641.
- Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy**, "Measuring technological innovation over the long run," *American Economic Review: Insights*, 2021, 3 (3), 303–320.
- Khan, Aubhik and Julia Thomas**, "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," *Econometrica*, 2008, 76 (2), 395–436.
- Koby, Yann and Christian K Wolf**, "Aggregation in Heterogeneous-Firm Models: A Sufficient Statistics Approach *," Technical Report 2020.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman**, "Technological innovation, resource allocation, and growth," *The quarterly journal of economics*, 2017, 132 (2), 665–712.
- Kormendi, Roger C and Philip G Meguire**, "Macroeconomic determinants of growth: cross-country evidence," *Journal of Monetary economics*, 1985, 16 (2), 141–163.
- Krusell, Per and Anthony A Smith Jr**, "Income and wealth heterogeneity in the macroeconomy," *Journal of political Economy*, 1998, 106 (5), 867–896.
- Kung, Howard and Lukas Schmid**, "Innovation, growth, and asset prices," *The Journal of Finance*, 2015, 70 (3), 1001–1037.
- Lee, Hanbaek**, "Striking While the Iron Is Cold: Fragility after a Surge of Lumpy Investments," *Journal of the European Economic Association*, 2026.

- Moscarini, Giuseppe and Fabien Postel-Vinay**, “The cyclical job ladder,” *Annual Review of Economics*, 2018, 10, 165–188.
- Ottonello, Pablo and Thomas Winberry**, “Financial heterogeneity and the investment channel of monetary policy,” *Econometrica*, 2020, 88 (6), 2473–2502.
- and —, “Capital, Ideas, and the Costs of Financial Frictions,” Working Paper 32056, National Bureau of Economic Research 2026.
- Ouyang, Min**, “On the Cyclicalities of R&D,” *Review of Economics and Statistics*, 2011, 93 (2), 542–553.
- Queralto, Albert**, “A model of slow recoveries from financial crises,” *Journal of Monetary Economics*, 2020, 114, 1–25.
- Ramey, Garey and Valerie A Ramey**, “Cross-country evidence on the link between volatility and growth,” *The American Economic Review*, 1995, 85 (5), 1138.
- Romer, Paul M**, “Endogenous technological change,” *Journal of political Economy*, 1990, 98 (5, Part 2), S71–S102.
- Schmitz, Tom**, “Endogenous growth, firm heterogeneity and the long-run impact of financial crises,” *European Economic Review*, 2021, 132, 103637.
- Sedláček, Petr**, “Lost Generations of Firms and Aggregate Labor Market Dynamics,” *Journal of Monetary Economics*, 2020, 111, 16–31.
- and **Vincent Sterk**, “The growth potential of startups over the business cycle,” *American Economic Review*, 2017, 107 (10), 3182–3210.
- Stadler, George W.**, “Business Cycle Models with Endogenous Technology,” *American Economic Review*, 1990, 80 (4), 763–778.
- Sterk, Vincent, Petr Sedláček, and Benjamin Pugsley**, “The Nature of Firm Growth,” *American Economic Review*, 2021, 111 (2), 547–579.
- Thomas, Julia K**, “Is lumpy investment relevant for the business cycle?,” *Journal of political Economy*, 2002, 110 (3), 508–534.
- Wicksell, Knut**, “Interest and Prices (English translation by RF Kahn, 1936),” 1898.
- Winberry, Thomas**, “Lumpy investment, business cycles, and stimulus policy,” *American Economic Review*, 2021, 111 (1), 364–96.

A Proofs of the Main Results

A.1 Proof of Proposition 1

Proof. Project values satisfy $\tilde{J}_t > 0$. Throughout we use that a raw idea, depreciated by one period's obsolescence and discounted to the present, is not expected to be worth more than it is today:

$$E_t[\Lambda_{t,t+1} \phi \tilde{J}_{t+1}] < \tilde{J}_t. \quad (\text{A.1})$$

When the value of a raw idea is constant, (A.1) reads $\beta\phi < 1$.

Let $\Delta_t(s) \equiv S_t^a(s) - S_t^n(s)$ denote the gain from launching. By (3.4) and (3.5),

$$\Delta_t(s) = -c + s\tilde{J}_t + E_t\Lambda_{t,t+1} \left[S_{t+1}(\eta_{i,t+1}) - S_{t+1}(\phi s + \eta_{i,t+1}) \right]. \quad (\text{A.2})$$

Step 1. We first bound how fast S_t can rise in s . Fix $x' > x$ and consider two innovators holding these stocks who face the same realizations of the aggregate and idiosyncratic shocks. Let τ be any launch date measurable with respect to the common information. If both innovators launch at τ , their stocks coincide from $\tau + 1$ onward, because each resets and re-enters with the same $\eta_{i,\tau+1}$; until then their stocks differ by $\phi^k(x' - x)$ after k periods of waiting. The payoffs from following τ therefore differ by exactly $E_t[\Lambda_{t,\tau}\phi^{\tau-t}\tilde{J}_\tau](x' - x)$.

Let τ' attain $S_t(x')$. It is also available to the innovator holding x , so

$$0 < S_t(x') - S_t(x) \leq E_t[\Lambda_{t,\tau'}\phi^{\tau'-t}\tilde{J}_{\tau'}](x' - x) \leq \tilde{J}_t(x' - x), \quad (\text{A.3})$$

the first inequality following from the same comparison at the launch date that attains $S_t(x)$ together with $\tilde{J} > 0$. The last inequality holds because (A.1) makes $\Lambda_{t,\tau}\phi^{\tau-t}\tilde{J}_\tau$ a supermartingale in τ , so its expectation is largest at $\tau = t$. In particular S_t is Lipschitz, hence continuous.

Step 2. Δ_t is strictly increasing. For $s' > s$, (A.2) and (A.3) applied one period ahead give

$$\begin{aligned} \Delta_t(s') - \Delta_t(s) &= (s' - s)\tilde{J}_t - E_t\Lambda_{t,t+1} \left[S_{t+1}(\phi s' + \eta_{i,t+1}) - S_{t+1}(\phi s + \eta_{i,t+1}) \right] \\ &\geq (s' - s) \left[\tilde{J}_t - \phi E_t[\Lambda_{t,t+1}\tilde{J}_{t+1}] \right] > 0, \end{aligned} \quad (\text{A.4})$$

where the last inequality is (A.1).

Step 3. Launching costs c and gains nothing at $s = 0$, so $\Delta_t(0) = -c < 0$. Setting $s = 0$ in (A.4) gives $\Delta_t(s) \geq -c + s[\tilde{f}_t - \phi E_t \Lambda_{t,t+1} \tilde{f}_{t+1}]$, which diverges. Since Δ_t is continuous and strictly increasing it has a unique zero $\bar{s}_t > 0$, and launching is optimal if and only if $s > \bar{s}_t$. $\square$

A.2 Proof of Lemma 1

Proof. Write $h = \bar{s}^*$, and let f^* be the associated stationary raw-stock density. Let Φ and φ denote the standard-normal distribution and density functions, and let ψ and Ψ denote the inflow density and distribution function. Extend both by zero on negative arguments and write $Q = 1 - \Psi$. With $Z = \Phi(\mu/\tau)$, the inflow median is

$$q_{1/2} = \mu + \tau \Phi^{-1}(1 - Z/2) > \max\{\mu, 0\}. \quad (\text{A.5})$$

Positivity is immediate from the support of the truncated normal; if $\mu > 0$, then $\Psi(\mu) = 1 - 1/(2Z) < 1/2$.

At the no-shock launch threshold, indifference gives

$$h\tilde{f}^* = c + \mathbb{E}\Lambda[S^*(\phi h + \eta) - S^*(\eta)] \geq c,$$

because the value function is nondecreasing. Hence $h \geq c/\tilde{f}^*$. Define

$$c_A^* = \frac{\tilde{f}^* q_{1/2}}{1 - \phi}. \quad (\text{A.6})$$

If $c > c_A^*$, then

$$u := (1 - \phi)h > q_{1/2}. \quad (\text{A.7})$$

For completeness, let $P_h = \int_h^\infty f^*(x) dx$. The stationary density satisfies

$$f^*(s) = P_h \psi(s) + \int_0^h f^*(x) \psi(s - \phi x) dx. \quad (\text{A.8})$$

Equivalently, with the stationary carry law

$$\rho_h = P_h \delta_0 + \phi^{-1} f^*(y/\phi) \mathbf{1}_{\{0 < y < \phi h\}} dy, \quad f^*(s) = \int \psi(s - y) \rho_h(dy).$$

The carry kernel minorizes $Q(h)\delta_0$, so its invariant law is unique. The display also gives a positive continuous raw density and $P_h \geq Q(h) > 0$.

For $z > 0$ and $d > 0$, define

$$J(z, d) = 2\Psi(z) - \Psi(z - d) - \Psi(z + d),$$

where Ψ is extended by zero on negative arguments. Fubini's theorem gives, for every feasible equal threshold movement $0 < \Delta \leq h$,

$$\int_{h-\Delta}^h f^*(x) dx - \int_h^{h+\Delta} f^*(x) dx = \int J(h - y, \Delta) \rho_h(dy). \quad (\text{A.9})$$

For every $y < \phi h$, put $z = h - y > u > q_{1/2} > \mu$. If $0 < \Delta \leq z$, then for $0 < v \leq \Delta$,

$$\frac{\psi(z - v)}{\psi(z + v)} = \exp\left(\frac{2v(z - \mu)}{\tau^2}\right) > 1,$$

so $J(z, \Delta) > 0$. If $\Delta > z$, then $J(z, \Delta) = 2\Psi(z) - \Psi(z + \Delta) > 0$, because $\Psi(z) > 1/2$. Every component in (A.9) is therefore strictly positive, which proves the lemma. $\square$

Remark. The two integrals are the masses whose launch decisions change. They also equal the respective total variation distances of post-decision carry from baseline: the affected mass moves between zero and strictly positive carry. Adding the next inflow contracts these distances but need not preserve their ordering. Equal threshold movements do not mean equal primitive shocks.

A.3 Proof of Proposition 2

Proof. Write $h = \bar{s}^*$, $a = \sigma$, $L = h(1 - a)$, $U = h(1 + a)$, and $\alpha = 1 - \phi^2$. We show that the cycle-average density at h is below its steady-state value for every $a \in (0, 1)$ once h exceeds a common finite bound. Let ψ and Ψ denote the inflow density and distribution, extended by zero on negative arguments, and put $Q = 1 - \Psi$. The transition kernel for incoming carry is

$$C_k z(y) = \int_0^{(k-y)_+} \psi(r) z(\phi(y + r)) dr + Q(k - y) z(0).$$

Let $\rho_+ = \rho_- C_L$, $\rho_- = \rho_+ C_U$, and $\rho_h = \rho_h C_h$. Lemma D.1 in Online Appendix D.2 establishes their existence and uniqueness. Their positive raw densities are $f_{\pm}(s) =$

$\rho_{\pm}[\psi(s - \cdot)]$ and $f_h(s) = \rho_h[\psi(s - \cdot)]$. Thus f_- faces L and f_+ faces U . Set $F = (f_- + f_+)/2$ and $D(h, a) = F(h) - f_h(h)$.

The density comparison reduces to the stocks whose decisions change. For $p_s(y) = \psi(s - y)$, Lemma D.1 gives the bounded solution $(I - C_h)\Gamma = p_h - f_h(h)$, normalized by $\Gamma(0) = 0$, on $[0, 2\phi h]$. Put $G(x) = \Gamma(\phi x)$ and $\bar{\rho} = (\rho_- + \rho_+)/2$. The phase equations imply

$$2\bar{\rho}(I - C_h) = \rho_-(C_L - C_h) + \rho_+(C_U - C_h).$$

Applying this identity to Γ and using $(C_k - C_h)\Gamma(y) = \int_h^k \psi(x - y)G(x) dx$ yields

$$2D(h, a) = \int_h^U f_+(x)G(x) dx - \int_L^h f_-(x)G(x) dx.$$

We compare these terms in three amplitude regions. Throughout, constants are independent of h and a ; auxiliary choices depend only on the fixed parameters (μ, τ, ϕ) .

First, small amplitudes. Choose $a_s \in (0, (1 - \phi)/2)$ satisfying (D.9), and write $\ell = 1 - a_s$, $v = 1 + a_s$, $d = \ell - \phi v$. In particular, $d > \phi(1 - \phi\ell)$. For $0 < a \leq a_s$, set $C_A = (C_L + C_U)/2$ and $\delta\rho = (\rho_- - \rho_+)/2$. Then $\delta\rho(I + C_A) = \bar{\rho}(C_U - C_L)/2$. Lemma D.2 gives a normalized solution $(I + C_A)\chi_s = p_s - b_s$ with $\chi_s(\phi x) > 0$ for $s, x \in [L, U]$. Consequently,

$$f_-(s) - f_+(s) = \int_L^U F(x)\chi_s(\phi x) dx > 0. \quad (\text{A.10})$$

The same lemma gives, uniformly on $[L, U]$,

$$G > 0, \quad \frac{G'}{G} = (1 + o(1)) \frac{\phi(h - \phi x - \mu)}{\tau^2}.$$

Incoming carry is bounded by $B = \phi v h$. Differentiating its density mixture gives $F'/F \leq -(dh - \mu)/\tau^2$. Since $h - \phi x \leq (1 - \phi\ell)h$, the strict gap above implies $(FG)' < 0$ for all sufficiently large h , uniformly over these amplitudes. With $\varepsilon = ha > 0$, the response identity and (A.10) imply

$$2D(h, a) \leq \int_0^\varepsilon \{F(h+t)G(h+t) - F(h-t)G(h-t)\} dt < 0.$$

This provides a finite bound H_S uniform over $0 < a \leq a_s$.

Second, amplitudes with $a \geq a_s$ and $L/h \geq r_0$. Set $r_0 = (\phi^3 + \phi^2)/2$ and $q = 1 - a_s/2$, so $g_0 = 2(1 - \phi) - \{\alpha + (q - \phi)^2\} > 0$. Lemma D.3 gives $G \geq 0$ on $[0, h]$ and

$$\int_h^U f_+ G \leq C(1 + h^3)e^{-(1-\phi)h^2/\tau^2 + Ch}.$$

The path lower bound on an interval inside $[L, h] \cap [\phi h, qh]$ yields (D.16):

$$\int_L^h f_- G \geq c_1 h e^{-(1-\phi)h^2/\tau^2 + g_0 h^2/(4\tau^2) - Ch}.$$

The interval width and path length depend only on ϕ and a_s . The lower bound therefore dominates uniformly in this region. Thus $D(h, a) < 0$ for $h > H_M$.

Third, amplitudes with $a \geq a_s$ and $0 < L/h \leq r_0$. The Gaussian convolution estimate (D.17) gives

$$F(h) \leq C(1+h)e^{-(\alpha+k_2)h^2/(2\tau^2)+Ch}, \quad k_2 = \frac{(\phi^2 - r_0)^2}{1 + \phi^2} > 0.$$

Let $V_n = \sum_{j=0}^{n-1} \phi^{2j}$ and $M_n = \mu \sum_{j=0}^{n-1} \phi^j$. Fix $n \geq 3$ with $V_n^{-1} < \alpha + k_2/2$, which is possible because $V_n^{-1} \rightarrow \alpha$. Lemma D.3 gives $f_h(h) \geq c_n e^{-(h-M_n)^2/(2\tau^2 V_n)}$. Hence

$$\frac{F(h)}{f_h(h)} \leq C(1+h)e^{-k_2 h^2/(4\tau^2)+Ch} \rightarrow 0$$

uniformly for $0 < L/h \leq r_0$, including sequences tending to zero. Thus $D(h, a) < 0$ for $h > H_L$.

These regions cover every $a \in (0, 1)$, including their boundaries. Using (A.6), choose

$$H = \max\{\tau, H_S, H_M, H_L\} < \infty, \quad c^* = \max\{c_A^*, \tilde{f}^* H\}. \quad (\text{A.11})$$

Lemma 1's proof gives $h \geq c/\tilde{f}^*$. Therefore $c > c^*$ implies both $h > H$ and that lemma's sufficient-lumpiness condition. With $a = \sigma$, $f_h = f^*$, and $F = f_\sigma$, the uniform inequality $D(h, a) < 0$ proves the proposition. The same finite cutoff applies to Lemma 1. $\square$

Online Appendix

The Distribution of Innovators and Medium-Term Business Cycles


Lixing Wang
Donghai Zhang

A. Additional Figures	2
B. Data Appendix	2
B.1. Sample Selection	2
B.2. R&D Adjustment: Discussion and Robustness Checks	3
B.3. Data Appendix for the Great Twins	8
C. Model Appendix	11
C.1. A Formal Representation of Compressing and Stretching	11
C.2. Launch Thresholds and Asymmetry	11
C.3. Great Moderation in Lumpy Capital Investment	14
C.4. Discussions of Model Assumptions	16
C.5. Equilibrium and the Solution Method	19
C.6. Adjustment Timing and Raw-Idea Stocks	22
C.7. The Great Twins When Ideas are Getting Harder to Find	22
D. Technical Proofs	23
D.1. The Proof of Lemma A.1	23
D.2. Auxiliary Results for Proposition 2	24
D.2.1. Common probability tools	25
D.2.2. Small amplitudes: a uniform derivative argument	26
D.2.3. Amplitudes bounded away from zero	28

A Additional Figures

We estimate aggregate demand shocks by fitting the model to U.S. GDP growth, as described in Section 5.3. Figure A.1 reports this in-sample fit and compares model-implied TFP growth with the utilization-adjusted TFP series, which is not used to estimate the shocks.

Figure A.1: Estimated Aggregate Shocks and External TFP Validation



Notes: The TFP series is from the Federal Reserve Bank of San Francisco and measures utilization-adjusted TFP following Fernald (2014). Because measured TFP contains movements beyond technological progress, including changes in misallocation, the red line reports its three-year moving average.

Interpreting the Evolution of TFP The fitted model provides a distributional interpretation of the evolution of TFP growth. The late-1990s dot-com boom raised R&D and drew down innovators' raw-idea stocks. The subsequent period of moderation shifted the innovation sector toward rebuilding these stocks, temporarily reducing the ideas assigned to new projects. This channel helps the model account for the subsequent path of TFP growth and complements explanations based on changes in information and communications technology (Cettè et al., 2016), technology diffusion (Anzoategui et al., 2019), and the difficulty of finding ideas (Bloom et al., 2020).

B Data Appendix

B.1 Sample Selection

We use annual observations from Compustat Fundamentals for 1968–2020, indexed by firm identifier (GVKEY) and fiscal year. We retain firms incorporated in the United States

that report in U.S. dollars. We exclude utilities (SIC 4900–4999), financial firms (SIC 6000–6999), non-operating establishments (SIC 9995), and industrial conglomerates (SIC 9997). An innovating firm is one that reports positive R&D expenditure at least once during the sample period. For the time-series analysis, we require at least five firm-year observations.

We remove observations with negative R&D expenditure and firm-years in which acquisitions exceed 5% of total assets. The acquisition screen prevents changes in reported R&D caused by the incorporation of an acquired firm’s research operations from being classified as internal adjustment. Table B.1 compares the industry composition of the resulting sample with the full Compustat population. The sample spans a broad set of industries but places greater weight on R&D-intensive manufacturing firms.

Table B.1: Industry Statistics

Industry	Number of Firms	Fraction	Fraction (full sample)
Agriculture, Forestry and Fishing	35	0.0049	0.0037
Construction	41	0.0057	0.0097
Mining	138	0.0191	0.0544
Retail Trade	104	0.0144	0.0502
Services	1480	0.2052	0.1585
Transportation & Public Utilities	159	0.0220	0.0745
Wholesale Trade	167	0.0231	0.0304
Manufacturing			
Chemicals and Allied Products	1025	0.1421	0.0717
Electronic and Other Electrical Equipment	932	0.1292	0.0472
Instruments and Related Products	841	0.1166	0.0413
Industrial Machinery and Equipment	852	0.1181	0.0404
Transportation Equipment	256	0.0355	0.0151
Other Manufacturing	1184	0.1641	0.1052

B.2 R&D Adjustment: Discussion and Robustness Checks

Section 2 documents that firm-level R&D adjustment is episodic and that large increases are procyclical. This subsection provides additional evidence on both findings. We first examine whether the episodic pattern is driven by tax incentives, firm scale, financing conditions, or the treatment of entry and exit. We then report additional evidence on the cyclicity of large adjustments.

R&D Tax Incentives The U.S. credit for increasing research activities is unlikely to induce the episodic pattern documented in the main text. The credit applies to qualified

research expenditures above a calculated base amount. This base depends on a firm's historical qualified research expenditures and is bounded below by 50% of current qualified expenditures. The tax schedule therefore rewards sustained research spending rather than the sharp intertemporal concentration of expenditures. Strategic timing of expenditures to obtain the credit is thus unlikely to account for the pronounced right tail in annual R&D changes.

Alternative Measures and Samples The episodic pattern also appears in innovation outcomes. Figure B.1a plots the distribution of annual changes in new patent counts for Compustat firms, using USPTO patent data matched to Compustat following [Arora et al. \(2021\)](#). Changes in patent counts display the same combination of many small adjustments and a pronounced right tail.


Firm Scale Firm growth does not account for the episodic pattern. Figure B.1b normalizes R&D expenditure by sales, while Figure B.1d restricts the sample to firms whose sales changed by less than 10% over the preceding three years. In both cases, annual R&D changes remain concentrated around small adjustments and retain a pronounced right tail. Figure B.1c further shows that the pattern persists among firms in the top two quintiles of average R&D expenditure. Episodic adjustment is therefore neither a mechanical consequence of changes in firm scale nor confined to firms with small R&D programs.

Financial Conditions Figure B.2 compares firms in the bottom and top quartiles of the debt-to-equity distribution. Both groups exhibit similar concentrations of small changes and pronounced right tails. The episodic pattern is therefore not concentrated among highly leveraged firms.


Cyclicity across Firm Size Figure B.3 compares the cyclicity of the spike rate across firms with different average R&D expenditures. Small-R&D firms are in the bottom two quintiles, while large-R&D firms are in the top two quintiles. The spike rate is procyclical in both groups, with correlations of 0.56 among small-R&D firms and 0.63 among large-R&D firms. The aggregate relationship is therefore not driven by firms with small R&D programs.

Pooled Spike Thresholds Table 1 reports the cyclicity of the spike rate separately for each threshold. Figure B.4 provides a complementary time-series view that pools spike definitions from 40% to 100%. The pooled spike rate comoves strongly with real GDP growth; the correlation is 0.52 with a p-value of 0.00.


Figure B.1: Robustness of Episodic R&D Adjustment




(a) Changes in the Number of New R&D Patents



(b) Changes in the R&D-to-Sale Ratio



(c) Changes in the R&D Expenditure (Large Firms)




(d) Firms with Stable Scale past 3 Years

Notes: Figure B.1a shows the histogram of changes in the number of new R&D patents annually, with data sourced from the USPTO. We select assignees with positive newly issued patents for at least five years. Figure B.1b depicts the histogram of changes in the R&D-to-sale ratio, excluding observations with non-positive sales. Figure B.1c plots the changes in R&D expenditure for large firms. In particular, firms are categorized into five groups based on their average R&D expenses, with small firms being in the first and second quintiles and large firms being in the fourth and fifth quintiles. Figure B.1d shows the distribution of changes in R&D expense for firms whose sales varied by less than 10% over the past three years.

Alternative Aggregate Indicators Because R&D expenditure is included in GDP, mechanical accounting could contribute to its measured comovement with GDP growth. Figure B.5 addresses this concern using GDP net of aggregate R&D expenditure and unemployment. The spike rate remains procyclical relative to GDP net of R&D and counter-cyclical relative to unemployment.

Entry, Exit, and Inaction Changes from zero to positive reported R&D may reflect entry into the sample or a change in accounting practice. The percent-change measure bounds these observations at 200%, and the right tail remains after they are excluded. Figure B.6a also recomputes the spike rate after excluding zero-to-positive changes; its procyclicality

Figure B.2: Episodic R&D Adjustment by Financial Leverage



Notes: Firms are categorized into low- and high-leverage groups using the debt-to-equity ratio. Low-leverage firms are in the first quartile, while high-leverage firms are in the fourth quartile.





Figure B.3: Cyclicity of R&D Spikes across Firm Size


Notes: Firms are assigned to quintiles based on average R&D expenditure. Small-R&D firms are in the first and second quintiles; large-R&D firms are in the fourth and fifth quintiles. Both the spike rate and real GDP are HP-filtered.

is preserved. Finally, Figure B.6b shows that the share of firms leaving R&D unchanged is countercyclical. This extensive-margin evidence complements the procyclicality of large increases.


Figure B.4: Cyclicalty of the Pooled-Threshold R&D Spike Rate



Notes: This figure plots the R&D spike rate and real GDP growth from 1984 to 2018. The R&D spike rate is the annual share of innovating firms making a spike adjustment, pooling thresholds from 40% to 100%. Changes below \$10,000 and observations in which acquisitions exceed 5% of total assets are excluded.



(a) GDP net of R&D Spending



(b) Unemployment

Figure B.5: Cyclicalty Using Alternative Aggregate Indicators

Notes: GDP net of R&D expenditure is calculated by subtracting aggregate R&D expenditure from real GDP. Real GDP, aggregate R&D expenditure, and unemployment are from FRED.




Figure B.6: Entry, Exit, and Inaction

Notes: The left panel plots the spike rate after excluding zero-to-positive R&D changes against GDP growth. The right panel plots the inactive rate against GDP. Following [Cooper and Haltiwanger \(2006\)](#), the inactive rate is the annual share of firms that leave R&D expenditure unchanged.

B.3 Data Appendix for the Great Twins

This subsection describes the measurement and robustness checks for the Great Twins evidence in Section 7.1. We first define recession depth, then examine alternative outcomes, historical recessions, R&D-intensity cutoffs, and financial controls.


Measuring the Depth of Recession We formally define the gap between the actual GDP and the pre-recession trend level as follows:

$$\text{Gap}_{t_0, t_1} = \log Y_{t_0} + g_m \times (t_1 - t_0) - \log Y_{t_1} \quad (\text{B.1})$$


where t_0 represents the initial period before the recession, and g_m represents the trend growth rate prior to the recession. Therefore, the term $\log Y_{t_0} + g_m \times (t_1 - t_0)$ calculates the expected log GDP at time t_1 , assuming the economy had continued to expand at the pre-recession trend rate from the initial time t_0 . Y_{t_1} denotes actual GDP at time t_1 .

We define recession depth as the average gap between actual GDP and its pre-recession trajectory:

$$\text{Depth of Recession}_{t_0} = \frac{1}{K} \sum_{k=10-K+1}^{10} \text{Gap}_{t_0, t_0+k} \quad (\text{B.2})$$



(a) Great Twins: TFP



(b) Great Twins: Beyond 2008

Figure B.7: The Great Twins Robustness Checks.

Notes: The left figure depicts the relationship between moderation and recession when the depth of the recession is defined using TFP. The cross-country TFP data is from the Penn World Database. The correlation between the degree of moderation and the depth of the recession is 0.37 ($p = 0.08$) for the OECD sample. The right figure presents the depth of recessions and the degree of moderation across time and countries. The correlation is 0.35 ($p = 0.00$). This pattern remains robust when the 2008 recession is excluded from the cross-country sample, with the correlation for the subsample excluding the 2008 recession standing at 0.34 ($p = 0.00$).

In practice, we select $K = 5$ because the recovery from the recession typically takes several years, during which the economy gradually aligns with its new trend. By taking the average, we further mitigate the fluctuations brought about by transitory shocks. Our empirical findings are robust to alternative choices of K .

The Great Twins Robustness Checks: TFP-Based Recession Depth Section 7.1 measures recession depth using real GDP. Because the model's medium-term output effects operate through technology, we repeat the cross-country exercise using TFP to measure recession depth. Figure B.7a shows that countries with lower pre-2008 volatility experienced larger TFP losses during the Great Recession, consistent with the model's mechanism.

The Great Twins Robustness Checks: Beyond the 2008 Recession In the main text, our evidence for the Great Twins focuses on the 2008 Recession. A potential concern is whether this characteristic is unique to the financial crisis of 2008. Figure B.7b extends this exercise to a cross-country and cross-historical recession dimension. We analyzed major recessions in 18 advanced countries starting from the end of WWII. Figure B.7b demonstrates that the Great Twins phenomenon exists across different countries and historical recessions. The positive correlation between pre-recession moderation and the loss of GDP remains significant, even when excluding the 2008 recession.

Table B.2: The Robustness of the New Empirics

	RD >1.	RD >1.5	RD >0.5	RD <1.	AC1	AC2	AC3	AC4
Degree of Moderation	9.38*** (2.43)	15.47*** (5.25)	6.73*** (2.43)	-1.08 (2.83)	12.70*** (3.26)	12.33** (4.41)	13.09*** (3.94)	13.07*** (4.37)
$\frac{(\text{Private Credit/GDP})_{2007}}{(\text{Private Credit/GDP})_{2000}}$							3.99 (2.33)	3.97 (2.56)
Real interest rate ₂₀₀₇						0.35 (0.43)		0.12 (0.45)
R-squared	0.38	0.35	0.18	0.00	0.41	0.36	0.43	0.44
R-squared Adj.	0.36	0.31	0.16	-0.02	0.38	0.28	0.36	0.33
const	2.39** (1.10)	0.17 (2.41)	2.86** (1.06)	4.94*** (1.23)	0.81 (1.48)	0.12 (2.14)	-4.23 (3.59)	-4.48 (3.98)
N	26	18	36	39	24	21	21	19

Notes: The data of real interest rate and private credit to GDP ratio are from the World Bank database. In particular, the ratio $\frac{(\text{Private Credit/GDP})_{2007}}{(\text{Private Credit/GDP})_{2000}}$ measures the aggregate leverage with respect to the baseline year of 2000. *** significant at the 1% level, ** significant at the 5% level, * significant at the 10% level.

The Great Twins Robustness Checks: Sample Selection We examine whether the Great Twins relationship is sensitive to the R&D-intensity cutoff. The first four columns of Table B.2 report results for countries with R&D expenditure above 1%, above 1.5%, above 0.5%, and below 1% of GDP, respectively. The moderation coefficient is positive and statistically significant in the three above-threshold samples, whereas it is statistically insignificant in the below-1% sample.

The Great Twins Robustness Checks: Additional Controls In the main text, we observe cross-country evidence suggesting a positive correlation between moderation and recession. An alternative interpretation of this evidence might be that the Great Moderation coincided with significant leveraging and a savings glut (Bagliano and Morana 2017), indirectly contributing to the recession. To test whether these financial factors could explain our findings, we propose the following regression:

$$Depth_i^{2008} = \beta_0 + \beta_1 Mod_i^{2007} + \beta X_i^{2007} + u_i$$

where X_i^{2007} additionally controls for the Real Interest Rate in 2007 and the ratio of Private Credit to GDP in 2007 compared to 2000. Table B.2 reports a specification without either financial control in column AC1, adds the real interest rate in AC2 and the credit ratio in AC3, and includes both in AC4. The moderation coefficient remains positive and statistically significant across these specifications.

C Model Appendix

C.1 A Formal Representation of Compressing and Stretching

This subsection formalizes the compressing and stretching effects described in Section 3. Starting from the stationary density f^* at a constant launch threshold $\bar{s}^*$, we compare next-period densities after equal threshold movements in either direction.

Lemma A.1. *Under the maintained truncated-normal inflows, consider a shock that shifts the launch threshold left- or rightward by Δ units from the steady-state value $\bar{s}^*$, with $0 < \Delta \leq \bar{s}^*$, and let $f_{-\Delta}$ and $f_{+\Delta}$ denote the density of raw-idea stocks in the following period under the $-\Delta$ and the $+\Delta$ shock, respectively.*

- (i) **(Compressing Effect)** *In response to a $-\Delta$ shock, which shifts $\bar{s}^*$ leftward, there exists $s^{-\Delta} > 0$ such that*

$$f_{-\Delta}(s) < f^*(s) \quad \text{for every } s > s^{-\Delta}. \quad (\text{C.1})$$

- (ii) **(Stretching Effect)** *If the shock shifts $\bar{s}^*$ rightward, there exists $s^{+\Delta} > 0$ such that*

$$f_{+\Delta}(s) > f^*(s) \quad \text{for every } s > s^{+\Delta}. \quad (\text{C.2})$$

In particular, if the launch threshold at $t + 1$ lies above $s^{-\Delta}$, then after a favorable shock fewer innovators stand at the margin than in the steady state; if instead it lies above $s^{+\Delta}$, after an adverse shock more of them do.

Proof. See Online Appendix D. □

C.2 Launch Thresholds and Asymmetry

Does a higher launch threshold make the launch response more asymmetric? At each initial threshold, we start from the stationary distribution generated by that threshold. We then compare the additional launches induced by a given percentage reduction with the launches postponed by an equal percentage increase. We first solve two examples without obsolescence, $\phi = 1$, and then study $\phi < 1$ numerically.

Write $h = \bar{s} > 0$, let f_h denote its stationary raw-idea stock density, and fix $0 < \varepsilon < 1$. A reduction from h to $h(1 - \varepsilon)$ induces the innovators between these two thresholds to

launch; denote their mass by L_h . An increase from h to $h(1 + \varepsilon)$ postpones launches by innovators in the corresponding interval above h ; denote their mass by R_h . Both responses are evaluated using the same pre-shock density f_h . We measure asymmetry as the difference between these masses divided by their sum:

$$\text{ASY}_\varepsilon(h) \equiv \frac{L_h - R_h}{L_h + R_h}, \quad L_h = \int_{h(1-\varepsilon)}^h f_h(x) dx, \quad R_h = \int_h^{h(1+\varepsilon)} f_h(x) dx. \quad (\text{C.3})$$

A positive value means that additional launches exceed postponed launches; zero indicates equal responses, and a negative value reverses the comparison. Across initial thresholds, we hold the inflow process and ε fixed and recompute f_h . Thus the percentage movement is fixed, while its absolute size, εh , rises with the initial threshold.

For $\phi = 1$, the stationary density satisfies

$$f_h(s) = P_h \psi(s) + \int_0^{\min\{s, h\}} f_h(x) \psi(s - x) dx, \quad P_h = \int_h^\infty f_h(x) dx, \quad (\text{C.4})$$

where P_h is the launch-and-reset probability and ψ is the idea-inflow density. The first term represents innovators who launched and reset in the previous period, so their current stocks consist of new inflows. The second represents innovators who carried stocks below h forward and added new inflows.

Exponential Distribution Let $\psi(s) = e^{-s}$ for $s \geq 0$. Substitution into Equation (C.4) and normalization give

$$f_h(s) = \begin{cases} P_h, & 0 \leq s \leq h, \\ P_h e^{-(s-h)}, & s > h, \end{cases} \quad P_h = \frac{1}{1+h}. \quad (\text{C.5})$$

The density is flat below the launch threshold and declines exponentially above it. Equal-length intervals therefore contain more innovators below the threshold than above it. Writing $d = \varepsilon h$, the two interval masses are $L_h = P_h d$ and $R_h = P_h(1 - e^{-d})$. Hence

$$\text{ASY}_\varepsilon(h) = \frac{\varepsilon h - 1 + e^{-\varepsilon h}}{\varepsilon h + 1 - e^{-\varepsilon h}}. \quad (\text{C.6})$$

For fixed ε , this expression is strictly increasing in h : its derivative is

$$\frac{2\varepsilon[1 - (1+d)e^{-d}]}{(d+1 - e^{-d})^2} > 0,$$

since $e^d > 1 + d$ for $d > 0$. The increase comes from comparing wider intervals as h rises. The density remains flat to the left, whereas it falls farther to the right, so the relative excess mass on the left grows. For a fixed absolute movement $d < h$, the asymmetry measure is constant across initial thresholds in this example.

Hybrid Gamma Distribution The second example shows why the relationship need not be globally increasing. Consider the equal mixture of unit-rate exponential and gamma inflows with shape two:

$$\psi(s) = \frac{1}{2}(1+s)e^{-s}, \quad s \geq 0.$$

Write $q = e^{-3h/2}$. Substitution into Equation (C.4) and normalization give

$$f_h(s) = \begin{cases} P_h \left(\frac{2}{3} - \frac{1}{6}e^{-3s/2} \right), & 0 \leq s \leq h, \\ P_h e^{-(s-h)} \left[\frac{2}{3} - \frac{q}{6} + \left(\frac{1}{3} + \frac{q}{6} \right) (s-h) \right], & s > h, \end{cases} \quad (C.7)$$

$$P_h = \left[1 + \frac{2h}{3} - \frac{1-q}{9} \right]^{-1}.$$

For $d = \varepsilon h$, direct integration yields

$$\ell_h \equiv \frac{L_h}{P_h} = \frac{2d}{3} - \frac{q}{9} (e^{3d/2} - 1),$$

$$r_h \equiv \frac{R_h}{P_h} = 1 - \left[1 + \left(\frac{1}{3} + \frac{q}{6} \right) d \right] e^{-d}.$$

Thus the exact asymmetry measure is


$$\text{ASY}_\varepsilon(h) = \frac{\ell_h - r_h}{\ell_h + r_h}.$$

For fixed ε , its limiting behavior is

$$\text{ASY}_\varepsilon(h) = -\frac{\varepsilon h}{8} + O(h^2) \quad \text{as } h \downarrow 0, \quad \text{ASY}_\varepsilon(h) \longrightarrow 1 \quad \text{as } h \rightarrow \infty.$$

At sufficiently low thresholds, the asymmetry measure is negative: a threshold increase postpones more launches than an equal reduction induces. At sufficiently high thresholds, the mass in the interval below the threshold dominates the mass in the corresponding interval above it, and the measure approaches one. Positive asymmetry at high

Figure C.1: Asymmetry and the Launch Threshold



Notes: We set $\varepsilon = 0.1$. Curves are fifth-degree polynomial fits to simulated asymmetry for thresholds between 0.1 and 3, using 15 threshold values in panel (a) and 10 in panel (b). Each simulation uses 50,000 innovators and 1,000 periods, discarding the first 101 periods. Weibull inflows have density $kx^{k-1}e^{-x^k}$, $x > 0$, with $k = 0.5$. The series labeled Pareto uses the Pareto II (Lomax) density $k(1+x)^{-k-1}$, $x \geq 0$, with $k = 1$. Truncated-normal inflows are $N(0,1)$ draws conditional on being nonnegative; hybrid gamma inflows have density $\frac{1}{2}(1+x)e^{-x}$, $x \geq 0$.

thresholds can therefore coexist with a relationship that is not globally increasing.

Numerical Analysis We next allow unused ideas to become obsolete. Figure C.1 reports the same experiment for Weibull, Pareto, truncated-normal, and hybrid gamma inflows, with $\phi = 0.9$ in panel (a) and $\phi = 0.8$ in panel (b). For each initial threshold, we simulate its stationary distribution and compare the launch responses to a 10% reduction and a 10% increase. At higher thresholds, all four plotted curves are positive and rising: additional launches exceed postponed launches, and this relative imbalance grows with the initial threshold over the displayed range.

C.3 Great Moderation in Lumpy Capital Investment

This subsection examines the role of one-sided project launch in generating distributional hysteresis. We consider a capital-investment model in which firms can both invest and disinvest, allowing distributional effects at the two adjustment boundaries to offset each other.

Following [Baley and Blanco \(2021\)](#), we present a parsimonious partial equilibrium model of lumpy investment. Consider a firm engaged in production using capital, contending with capital adjustment frictions amidst a constant real interest rate r . The firm's

output, influenced by both aggregate and idiosyncratic productivity components, is described by:

$$y_{i,t} = (Z_t e_{i,t})^{1-\alpha} k_{i,t}^\alpha, \quad \text{where } \alpha < 1. \quad (\text{C.8})$$

Productivity components evolve according to:

$$z_{t+1} = \varepsilon_{t+1}^z, \quad \log e_{i,t+1} = \log e_{i,t} + \varepsilon_{i,t+1}^e. \quad (\text{C.9})$$

where z_t is the log representation of Z_t . Capital stock adjustments, either through acquisition or divestment, incur an adjustment cost $\theta_{i,t}$, proportional to total productivity:

$$\theta_{i,t} = Z_t e_{i,t} \Theta(i_t, \Delta N_t). \quad (\text{C.10})$$

where N_t is a Poisson counter with arrival rate λ :

$$\Theta(i_t, \Delta N_t) = \begin{cases} 0 & \text{if } i_t = 0 \text{ or } \Delta N_t = 1, \\ \theta^- & \text{if } i_t > 0 \text{ and } \Delta N_t = 0, \\ \theta^+ & \text{if } i_t < 0 \text{ and } \Delta N_t = 0. \end{cases} \quad (\text{C.11})$$

Assuming no capital depreciation, capital stock evolves according to:

$$k_{i,t} = k_{i,t-1} + i_{i,t}. \quad (\text{C.12})$$

With a discount rate β , the firm's sequential decision problem is to maximize its value by choosing optimal investments:

$$V(k_{i,t-1}; Z_t, e_{i,t}) = \max_{i_{i,t}} \left\{ (Z_t e_{i,t})^{1-\alpha} k_{i,t}^\alpha - r k_{i,t} - \theta_{i,t} + \beta E_t[V(k_{i,t}; Z_{t+1}, e_{i,t+1})] \right\}. \quad (\text{C.13})$$

The investment policy features an inactive region $\mathcal{R}$ defined by $(k^-(Z, e), k^+(Z, e))$ and an optimal capital level $k^*(Z, e)$, where no capital adjustment occurs within $\mathcal{R}$; outside this region:

$$i_{i,t} = k^*(Z_t, e_t) - k_{i,t-1}. \quad (\text{C.14})$$

Following [Baley and Blanco \(2021\)](#), the optimal policy is determined by $\{k^-, k^*, k^+\} = e \times \{Zx^-, Zx^*, Zx^+\}$. We define the capital gap as $\tilde{x}_{i,t} = \log(x^*) - \log(\frac{k_{i,t}}{e_{i,t}})$.

This formulation aligns the evolution of the capital gap with the dynamics of unused ideas in our lumpy R&D model, where the capital gap's dynamics are represented as:

$$\tilde{x}_{i,t} = (\tilde{x}_{i,t-1} + \varepsilon_{i,t}^e) \cdot I\left\{\log(x^*) - \log(x^+) - z_t \leq \tilde{x}_{i,t-1} + \varepsilon_{i,t}^e \leq \log(x^*) - \log(x^-) - z_t\right\}$$

Based on this parsimonious $\mathbb{S}$ model, Figure [C.3](#) illustrates why the mechanism emphasized in this paper is dampened in a conventional lumpy capital investment framework. The figure shows that a positive shock, which shifts the inactive region to the left, compresses the left tail of the capital gap distribution while stretching the right tail. A negative shock has a symmetric effect. As a result, the impact of aggregate shocks on the overall distribution is limited, due to the symmetric and offsetting nature of these adjustments in the $\mathbb{S}$ framework.


For the quantitative comparison, we allow persistent aggregate productivity shocks and different rates of costless adjustment opportunities for investment and disinvestment. We calibrate the lumpy investment model to Compustat capital investment data, as shown in Figure [C.2](#). A comparison of the average distribution of the capital gap with the steady-state distribution of the capital gap in Figure [C.2b](#) shows that the effects of aggregate fluctuations are negligible.

To conclude this discussion, it is also important to emphasize that even if certain assumptions lead to the average distribution of the capital gap having less mass around the extensive margin compared to the steady-state distribution, it remains uncertain whether this would result in the moderation-recession connection observed in the lumpy R&D model. The reason is that technological change exhibits greater persistence than fluctuations in capital. The dynamics of the aggregate level of capital may intertwine with shocks of business cycle frequency, implying that the lessened volatility of one, caused by the smaller volatility of another, would result in both neutralizing each other.


C.4 Discussions of Model Assumptions

Depreciation of ideas The obsolescence of ideas can be understood as an innovative idea gradually losing its business value, potentially because the intended market diminishes or is overtaken by emerging, more advanced technologies. In our closed-economy model, idea depreciation can reflect global competition, where innovators in other countries may

Figure C.2: Calibrated Lumpy Investment Model



(a) Investment Distribution



(b) Distribution of Capital Gap


Notes: We calibrate the model to match the quarterly distribution of capital investment in Compustat since 1986. In the calibrated model The exogenous adjustment rate for disinvestment, $\lambda^- = 0.05$ and the positive investment counterpart $\lambda^+ = 0.42$. The adjustment thresholds are asymmetric: $\log(x^+) - \log(x^*) = 0.3$, $\log(x^-) - \log(x^*) = -0.2$. The standard deviation of idiosyncratic productivity $\sigma^e = 0.26$. Figure C.2a compares the calibrated cross-sectional investment to its empirical counterpart. By fixing $\rho_z = 0.95$, we also calibrate σ_z to match a quarterly standard deviation of aggregate investment rate of 0.011. Figure C.2b compares the steady state and stationary distribution of capital gap.

have similar ideas and claim property rights, contributing to the idea depreciation in the home country.


Assumptions that Keep the Model Tractable In our model, there are two critical state variables. One is the distribution of $s_{i,t}$, which represents the amount of unused ideas held by innovators. It determines the expected growth of A_t and the response of A_t to demand shocks. The other is the distribution of $x_{i,t}$, which represents the amount of R&D ideas among innovators. It influences the demand for skilled labor. However, since distributions are infinite dimensional variables, employing numerical methods to handle them would significantly increase the model's complexity, and reduce its interpretability. Therefore, we attempt to maintain the model's tractability by restricting model parameters.

It is noted that the fluctuation of A_t primarily influences expectations of future consumption. In the Business Cycle block, the impact on future consumption expectations stems from the consumption Euler equation, which implies the intertemporal substitution of consumption. To negate the impact of the growth of A_t , we assume a monetary policy that respond to productivity.


Figure C.3: The Effect of Aggregate Shocks on the Distribution of Capital Gap




(a) Positive shocks shift inactive region to the left



(b) The distribution after a positive shock



(c) Negative shocks shift inactive region to the right



(d) The distribution after a negative shock

Regarding the Innovation block, the fluctuation of A_t mainly affects the Household's discounting of cash flows. First, regarding the value of each variety, V_t , we assume that each newly produced variety corresponds to a one-period right to claim the profit. This assumption ensures that the value of a variety is not affected by discounting of future cashflows. As for the remaining values J_t , $\tilde{J}_t$, and S_t , since they are proportional to A_t following general model assumptions, they are also not affected by the growth component of the discount factor.

The distribution of $x_{i,t}$ determines the total demand of skilled labor and hence impacts the wage of skilled labor. The wage of skilled labor enters the innovators' problem by affecting the cost of adoption. To eliminate the impact of the distribution of $x_{i,t}$ on $J_{i,t}$, we assume that the supply of skilled labor is perfectly elastic. This is achieved by setting $\eta = 0$.

Lastly, let us summarize the transmission mechanisms of the model under the model specifications discussed earlier. Apart from N_t , whose magnitude is determined by the magnitude of on-going research, and R_t , which includes a component linked to growth

expectations, all equilibrium values of variables in the Business Cycle Block are solely determined by the state variable q_t . This is consistent in the General Equilibrium. The Business Cycle Block exerts influence on the Innovation Block by affecting the value of IPR, the cost of idea adoption, and discounting; this, in turn, influences the distribution of idea stock and impacts current and future growth. On the one hand, however, within the Business Cycle Block, only households adjust consumption based on expectations of future income growth, a state determined in the Innovation Block. Central bank monetary policy accounts for these expectations and neutralizes changes in future income growth expectations within the household's consumption Euler equation by adjusting the nominal interest rate. On the other hand, though the magnitude of N_t is determined by the amount of on-going research, X_t , due to the perfectly elastic supply of skilled labor, the equilibrium quantity of skilled labor does not influence other Business Cycle Block variables. In a full-fledged Heterogeneous Agent model where distribution matters, household consumption and labor decisions are generally influenced by the distribution. In our model, the Business Cycle Block transmits the effects of preference shocks through consumption and labor choices to the Innovation Block. However, state variables determined within the Innovation Block do not reciprocally affect consumption and labor choices. This feature significantly simplifies the model's derivation and ensures its overall tractability, while preserving the key mechanism highlighted in the model.

Passive Adoption and Time to Build The transformation of R&D ideas into varieties undergoes two stages: project building and adoption. At the aggregate level, these two stages combined determine the persistence of R&D. At the micro-level, the time-to-build assumption primarily aims to match the persistence of the spike rate, meaning that firms' R&D spikes can occur during the boom period and the subsequent periods. In [Comin and Gertler \(2006\)](#), adoption is endogenous. We treat it as exogenous mainly because in both these settings, adoption occurs after the R&D decision is made, and therefore, they do not affect the mechanism described in the main text. This approach simplifies the model by focusing on the primary dynamics of R&D decision-making, without delving into the complexities of the adoption process, which, while important, do not directly influence the core mechanism under investigation.

C.5 Equilibrium and the Solution Method

This subsection defines equilibrium and describes the model's sequential solution. We first solve the business-cycle block, then innovators' decision rules, and finally the evolu-

tion of the innovation distributions and aggregate technology.

Equilibrium A recursive competitive equilibrium in this model is a set of value functions $\{S_{i,t}(s), \tilde{J}_t, J_t, V_t\}$, adjustment threshold $\bar{s}_t$, quantities $\{C_t, Y_t, Y_{t,m}, L_t, N_t, M_t, A_t\}$, prices $\{W_t, W_t^{RD}, \Pi_t, P_{t,m}, P_t^M, R_t\}$, and distributions of $s_{i,t}$ and $x_{i,t,j}$ such that all agents in the economy behave optimally, the distributions of firms and projects are consistent with decision rules, and all markets clear:

1. Intermediate Good Producer: Taking prices as given, the demand of L_t and M_t solves the intermediate good producer's problem.
2. Intermediate Input Producer: Taking other prices as given, the intermediate input producer optimally sets the price that maximizes the intermediate input producer's profits.
3. Innovator: Taking all prices as given, $V_t, J_t, \tilde{J}_t$ solves the corresponding Bellman equations. $S_{i,t}$ solves the Bellman equation defined by the innovator's choice to build a new R&D project, which yields the associated adjustment threshold $\bar{s}_t$.
4. Households: Taking prices as given, C_t and the supply of L_t and N_t satisfy the household's optimality conditions.
5. Retailer: The retailer maximizes profits such that the New Keynesian Phillips Curve holds.
6. Central Bank: The Taylor rule holds.
7. The distribution of $s_{i,t}$ and $x_{i,t,j}$ evolves according to the adjustment threshold $\bar{s}_t$.
8. All markets (final good, intermediate good, intermediate input, labor, skilled labor) clear.

Solution Method: The Business Cycle Block Given the separability of the variables within the New Keynesian (NK) and Production block from those in the Innovation Block, we employ the canonical Blanchard-Kahn method for solving the linearized system of equations to a first-order approximation. The linearized system of equations is presented in the following.

The Linearized Equations in the NK and Production Block Define the normalized variables $\tilde{C}_t \equiv C_t/A_t$, $\tilde{Y}_{t,m} \equiv Y_{t,m}/A_t$, $\tilde{M}_t \equiv M_t/A_t^\theta$, $\tilde{W}_t \equiv W_t/(A_t P_t)$, $\tilde{W}_t^{RD} \equiv W_t^{RD}/(A_t P_t)$, and $\tilde{P}_{t,m} \equiv P_{t,m}/P_t$. Lowercase variables denote their logs, not deviations from steady state. Let $g_t \equiv \log(A_t/A_{t-1})$, $i_t \equiv \log R_t$, $\rho \equiv -\log \beta$, and $\pi_t \equiv \Pi_t - 1$.

Household's supply of labor:

$$\tilde{c}_t + \varphi l_t = \tilde{w}_t, \quad \tilde{c}_t + \eta n_t = \tilde{w}_t^{RD} \quad (\text{C.15})$$

Household's consumption Euler equation:

$$\tilde{c}_t = E_t \tilde{c}_{t+1} - (i_t - E_t \pi_{t+1} - E_t g_{t+1} + E_t \Delta \log q_{t+1} - \rho) \quad (\text{C.16})$$

The Intermediate goods producer's demand for M_t :

$$\log(\gamma) + (1 - \gamma)l_t - (1 - \gamma)\tilde{m}_t = \log \vartheta \quad (\text{C.17})$$

The Intermediate goods producer's demand for L_t :

$$\log(1 - \gamma) - \gamma l_t + \gamma \tilde{m}_t = \tilde{w}_t - \tilde{p}_{t,m}. \quad (\text{C.18})$$

Retailer's price-setting problem (NKPC):

$$\pi_t = \beta E_t \pi_{t+1} + \frac{1}{\kappa(\mu - 1)} (\tilde{p}_{t,m} + \log \mu). \quad (\text{C.19})$$

Monetary policy:

$$i_t = \rho + g_{t+1} + \varphi_\pi \pi_t \quad (\text{C.20})$$

The intermediate-good producer's first-order condition and the markup $P_{t,j}^M = \vartheta P_{t,m}$ imply $\int_0^{A_t} M_{t,j} dj = (\gamma/\vartheta)Y_{t,m}$. Goods-market clearing then gives $C_t = (1 - \gamma/\vartheta)Y_{t,m}$. Using the normalizations above, its logarithmic form is

$$\tilde{c}_t = \log \left(1 - \frac{\gamma}{\vartheta} \right) + \tilde{y}_{t,m}. \quad (\text{C.21})$$

Solution Method: Innovation Block The restrictions described in Online Appendix C.4 allow us to solve innovators' decision rules without including the cross-sectional idea-stock distributions as state variables in their Bellman equations. We solve these equations by value function iteration, using the skilled-labor wage w_t^{RD} and the growth-adjusted stochastic discount factor $\tilde{\Lambda}_{t,t+1}$ from the NK and Production Block.

We then propagate the raw-idea stock distribution using Equation (3.6) and the en-

ogenous idea inflows in Section 4.1. Aggregating the ideas assigned to new projects and applying Equations (4.7)–(4.10) gives the stocks of pending and built R&D ideas and technology growth.

C.6 Adjustment Timing and Raw-Idea Stocks

Raw-idea stocks are not directly observable. This subsection explains how the distribution of time elapsed since firms' last R&D spike helps evaluate the model's implied distribution of these stocks.

Let $TE_{i,t}$ denote the time elapsed since innovator i 's last R&D spike. The following lemma links the distribution of this observable to the distribution of raw-idea stocks.

Lemma C.1. *Let $\mathcal{G}_t$ denote the cumulative distribution function of $s_{i,t}$, let $\mathcal{T}_t(n) = \Pr(TE_{i,t} = n)$ denote the probability mass function of $TE_{i,t}$, and let $F_t(s | n) = \Pr(s_{i,t} \leq s | TE_{i,t} = n)$ denote the conditional cumulative distribution function of $s_{i,t}$. Then*

$$\mathcal{G}_t(s) = \sum_{n \in \mathcal{N}} F_t(s | n) \mathcal{T}_t(n). \quad (\text{C.22})$$

The result follows directly from the law of total probability.

Elapsed-time frequencies provide the weights assigned to different groups of innovators, while the model determines the distribution of raw-idea stocks within each group.


C.7 The Great Twins When Ideas are Getting Harder to Find

Recent empirical evidence (Bloom et al., 2020; Kogan et al., 2017) suggests that ideas have been getting harder to find over the past four decades. In our model, this phenomenon is captured by a decline in innovation productivity, represented by a lower value of ζ . To assess how this trend affects the impact of the Great Twins, we vary ζ and simulate the model using the same series of estimated aggregate shocks.

Because lower innovation productivity implies a lower trend growth rate, we evaluate the excess loss of growth. This is the number of years of trend growth lost due to the 2008–2009 shocks, conditional on the 2008 macroeconomic state, relative to the average loss across all states.

When ζ falls by 10% relative to the baseline, the excess loss of growth from the Great Recession increases from 0.8 years to 1.4 years, a 75% rise. This indicates a more pronounced and persistent scarring effect when ideas become harder to find.

Figure C.4: The Great Twins when Ideas are Harder to Find



D Technical Proofs

D.1 The Proof of Lemma A.1

Proof. Write $h = \bar{s}^*$. Starting from the steady state, a threshold moved to $\bar{s}^* \mp \Delta$ for one period leaves the density

$$f_{\mp\Delta}(s) = \psi(s) \int_{\bar{s}^* \mp \Delta}^{\infty} f^*(x) dx + \int_0^{\bar{s}^* \mp \Delta} f^*(x) \psi(s - \phi x) dx, \quad (\text{D.1})$$

the first term collecting the innovators who launch and reset to a fresh inflow and the integral those who do not. Subtracting (A.8), the two terms combine over the range of stocks whose behavior the shock changes:

$$\begin{aligned} f_{-\Delta}(s) - f^*(s) &= - \int_{\bar{s}^* - \Delta}^{\bar{s}^*} f^*(x) [\psi(s - \phi x) - \psi(s)] dx, \\ f_{+\Delta}(s) - f^*(s) &= + \int_{\bar{s}^*}^{\bar{s}^* + \Delta} f^*(x) [\psi(s - \phi x) - \psi(s)] dx. \end{aligned} \quad (\text{D.2})$$

Everything turns on the bracket, which compares the two routes to a stock of s : an innovator who carried ϕx forward needs an inflow of $s - \phi x$, while one who launched and reset needs s . For the truncated normal, when $s \geq \phi x$,

$$\frac{\psi(s - \phi x)}{\psi(s)} = \exp \left[\frac{\phi x (s - \mu - \phi x / 2)}{\tau^2} \right].$$

The bracket is therefore positive if $s > \mu + \phi x / 2$; the first route must also be available. If $s < \phi x$, the innovator who carried ϕx forward has already passed s , $\psi(s - \phi x) = 0$, and

the bracket is $-\psi(s) < 0$.

It therefore suffices to keep s above every stock the shock leaves carried forward and above μ plus half that stock. In (D.2) the variable x ranges over $[\bar{s}^* - \Delta, \bar{s}^*]$ in the first line and over $[\bar{s}^*, \bar{s}^* + \Delta]$ in the second, so ϕx does not exceed $\phi \bar{s}^*$ and $\phi(\bar{s}^* + \Delta)$ respectively. Take

$$\begin{aligned} s^{-\Delta} &= \max\{\phi \bar{s}^*, \mu + \phi \bar{s}^*/2\}, \\ s^{+\Delta} &= \max\{\phi(\bar{s}^* + \Delta), \mu + \phi(\bar{s}^* + \Delta)/2\}. \end{aligned} \quad (\text{D.3})$$

For $s > s^{-\Delta}$ every x in the first range satisfies $s > \phi x$ and $s > \mu + \phi x/2$, and the bracket is strictly positive throughout, so the integral is strictly positive and (C.1) follows from the minus sign. For $s > s^{+\Delta}$ the same argument applies to the second range, and (C.2) follows from the plus sign. Both integrals are taken against a strictly positive weight: $f^* > 0$ on $[0, \bar{s}^*)$ by the positivity step in the proof of Lemma 1, and $f^* \geq P_h \psi > 0$ beyond $\bar{s}^*$ by (A.8). $\square$

Remark. The two cutoffs (D.3) account both for the largest carry and for the location of the inflow mode. Those who launch earlier would have carried at most $\phi \bar{s}^*$, and those who launch later carry at most $\phi(\bar{s}^* + \Delta)$. Above the corresponding cutoff the comparison is unambiguous; below it a carried stock can overshoot s , and the sign can reverse. Both comparisons hold at the baseline launch threshold if $\bar{s}^* \geq \mu + \phi \bar{s}^*/2$, $\phi \Delta \leq (1 - \phi) \bar{s}^*$, and $\bar{s}^* \geq \mu + \phi(\bar{s}^* + \Delta)/2$. Endpoint equalities are allowed because the upper endpoint of each affected stock interval has zero probability. At $\phi = 1$ the condition cannot be met, and the stretching comparison in fact reverses at $\bar{s}^*$: obsolescence is what places the effect at the margin.

D.2 Auxiliary Results for Proposition 2

This section establishes the technical bounds used in the proof in Appendix A.3. The bounds must hold uniformly over the fluctuation amplitude to yield one setup-cost cutoff. We therefore treat small amplitudes and amplitudes bounded away from zero separately.

Write $h = \bar{s}^*$ and $a = \sigma$, keeping τ for the standard deviation before truncation. Let

$$\begin{aligned} \psi(x) &= c_0 e^{-(x-\mu)^2/(2\tau^2)} \mathbf{1}_{\{x \geq 0\}}, & Z &= \Phi(\mu/\tau), \\ c_0 &= (\tau \sqrt{2\pi} Z)^{-1}, & \alpha &= 1 - \phi^2. \end{aligned}$$

and let Ψ be the inflow distribution function, extended by zero on negative arguments, with $Q = 1 - \Psi$. Throughout this proof, constants in estimates may change and depend only on (μ, τ, ϕ) and explicitly fixed auxiliary constants, never on the varying threshold or amplitude. Let $L = h(1 - a)$ and $U = h(1 + a)$.

D.2.1 Common probability tools

Let Y_t be incoming carry, so $s_t = Y_t + \eta_t$. On bounded measurable functions the exact decision kernel is

$$C_k z(y) = \int_0^{(k-y)_+} \psi(r) z(\phi(y+r)) dr + Q(k-y) z(0). \quad (\text{D.4})$$

In particular, $C_k z(y) = z(0)$ when $y \geq k$. For $L = h(1 - a)$, $U = h(1 + a)$, write

$$\begin{aligned} \rho_+ &= \rho_- C_L, & \rho_- &= \rho_+ C_U, & \rho_h &= \rho_h C_h, \\ f_{\pm}(s) &= \rho_{\pm}[\psi(s - \cdot)], & f_h(s) &= \rho_h[\psi(s - \cdot)]. \end{aligned}$$

Thus f_- faces the low threshold and f_+ the high threshold. Put $F = (f_- + f_+)/2$ and $D(h, a) = F(h) - f_h(h)$.

Lemma D.1 (Invariant laws, Gaussian bounds, and response identity). *For every $h > 0$ and $0 \leq a \leq 1$, these carry laws exist uniquely; their raw densities are bounded, positive and continuous on $[0, \infty)$ (right continuous at zero). Define*

$$e_\eta = \mu + \tau \phi(\mu/\tau)/Z, \quad M_c = \phi e_\eta / (1 - \phi), \quad M_s = e_\eta / (1 - \phi), \quad B_0 = \mu + M_c.$$

Each phase, including the baseline, satisfies

$$\begin{aligned} f(x) &\leq c_0 e^{-\alpha(x-B_0)_+^2/(2\tau^2)}, & \mathbb{E}Y &\leq M_c, \\ \Pr(s \geq x) &\leq e^{-\alpha(x-M_s)_+^2/(2\tau^2)}. \end{aligned} \quad (\text{D.5})$$

Let $p_s(y) = \psi(s - y)$ and $m = f_h(h)$. On $[0, 2\phi h]$ there is a unique bounded normalized solution $(I - C_h)\Gamma = p_h - m$, $\Gamma(0) = 0$. For $G(x) = \Gamma(\phi x)$,

$$2D(h, a) = \int_h^U f_+(x) G(x) dx - \int_L^h f_-(x) G(x) dx. \quad (\text{D.6})$$

Under the baseline threshold rule, $G(x)$ is the cumulative change in density at h from carrying ϕx rather than resetting to zero. The response identity weights each set of changed launch decisions by this subsequent effect.

Proof. All kernels with $k \leq 2h$ map into $[0, 2\phi h]$ and minorize $Q(2h)\delta_0$. Their two-step products strictly contract probabilities in total variation, giving the unique phase pair

and convergence along each phase; the fixed-threshold case is identical. Each carry law has a positive atom at zero and an L^1 density elsewhere. Convolution with bounded ψ is continuous for the latter part by continuity of translations in L^1 ; the atom gives a positive multiple of ψ . This proves all density claims, including the baseline representation

$$\rho_h = P_h \delta_0 + \phi^{-1} f_h(y/\phi) \mathbf{1}_{\{0 < y < \phi h\}} dy, \quad f_h(s) = \int \psi(s-y) \rho_h(dy). \quad (\text{D.7})$$

The inflow log moment generating function is $K(t) = \mu t + \tau^2 t^2/2 + \log \Phi(\mu/\tau + \tau t) - \log Z$. Writing $w(z) = \varphi(z)/\Phi(z)$ gives $(\log \Phi)'' = -w(z)(z + w(z)) < 0$ (the second factor is a positive truncated-normal mean). Thus

$$\mathbb{E} e^{t\eta} \leq e^{e\eta t + \tau^2 t^2/2}, \quad t \geq 0. \quad (\text{D.8})$$

Coupling from zero with the unrestricted recursion bounds carry by $Y_\infty := \sum_{j \geq 1} \phi^j \eta_j$ and raw stock by $S_\infty := \eta_0 + Y_\infty$. Their exponential moments are bounded, respectively, by

$$e^{M_c t + \tau^2 \phi^2 t^2/(2\alpha)} \quad \text{and} \quad e^{M_s t + \tau^2 t^2/(2\alpha)}.$$

Chernoff's bound gives the tail in (D.5). For the density use $\psi(r) \leq c_0 e^{\mu t + \tau^2 t^2/2 - tr}$ and optimize after averaging over Y at $t = \alpha(x - B_0)_+/\tau^2$.

Oscillation contraction gives uniform convergence of $\sum_{j \geq 0} C_h^j(p_h - m)$; subtracting its value at zero gives Γ . No continuity at $y = h$ is needed. With signed integration limits,

$$(C_k - C_h)z(y) = \int_h^k \psi(x-y)[z(\phi x) - z(0)] dx.$$

For $\bar{\rho} = (\rho_- + \rho_+)/2$ the phase equations give $\bar{\rho}(I - C_h) = \{\rho_-(C_L - C_h) + \rho_+(C_U - C_h)\}/2$. Apply this to Γ and use Fubini on the bounded strips to obtain (D.6). $\square$

D.2.2 Small amplitudes: a uniform derivative argument

Choose $a_s \in (0, (1 - \phi)/2)$, depending only on ϕ , so small that, with

$$\ell = 1 - a_s, \quad v = 1 + a_s, \quad d = \ell - \phi v, \quad D_0 = v - \phi \ell, \quad J_0 = \frac{(\ell - \phi^2 v)^2}{1 + \phi^2},$$

$$d > 0, \quad \ell > D_0, \quad \min(J_0, 2d^2) > D_0^2, \quad d > \phi(1 - \phi \ell). \quad (\text{D.9})$$

At $a_s = 0$ all are strict; in particular $J_0 - D_0^2 = 2\phi(1 - \phi)^2/(1 + \phi^2)$ and $d - \phi(1 - \phi \ell) = (1 - \phi)^2$. Continuity proves the choice exists.

Lemma D.2 (Boundary resolvent). *Put $I_h = [\ell h, v h]$, $X_h = [0, B]$ with $B = \phi v h$, and $J_h = [\phi \ell h, B]$. For any threshold probability κ on I_h , write $C_\kappa = \int C_k \kappa(dk)$. For sufficiently large h ,*

any $s \in I_h$, and $\theta \in \{-1, 1\}$, there is a unique normalized $C^1(X_h)$ solution

$$z - \theta C_\kappa z = p_s - b_s, \quad z(0) = 0,$$

for some constant b_s . With $q_s = p'_s = (s - \cdot - \mu)p_s/\tau^2$,

$$\sup_{s \in I_h, y \in J_h, \kappa, \theta} (|z(y)/p_s(y) - 1| + |z'(y)/q_s(y) - 1|) \longrightarrow 0. \quad (\text{D.10})$$

Proof. Since $B < \ell h$, let A_κ be the integral (nonreset) part of C_κ on X_h . On C^1 functions vanishing at zero use the complete norm $\|z'\|_\infty$. The centered map $Tz = C_\kappa z - C_\kappa z(0)$ satisfies

$$(C_\kappa z)' = \phi A_\kappa z' - \int \psi(k - \cdot) z(\phi k) \kappa(dk), \quad \|(Tz)'\|_\infty \leq [\phi + B\psi(dh)] \|z'\|_\infty.$$

For large h its coefficient is at most $(1 + \phi)/2 < 1$, uniformly in κ . The Neumann series for $z = p_s - p_s(0) + \theta Tz$ therefore converges in this norm, proving existence, uniqueness, and

$$\|z'\|_\infty \leq 2\|q_s\|_\infty/(1 - \phi), \quad \|q_s\|_\infty \leq Che^{-d^2 h^2/(2\tau^2) + Ch}. \quad (\text{D.11})$$

Let $R_\theta = (I - \theta\phi A_\kappa)^{-1}$; its sup norm is at most $(1 - \phi)^{-1}$. To estimate $R_\theta q_s - q_s$, dominate each surviving path of A_κ^n by the unrestricted carry kernel P^n applied to $q_s^{\text{ext}}(y) = |s - y - \mu|\psi(s - y)/\tau^2 \geq 0$. Completion of the square gives, for $t \geq 0$,

$$|r - \mu|\psi(r)e^{tr}/\tau^2 \leq c_0(t + \tau^{-1})e^{\mu t + \tau^2 t^2/2}.$$

Using (D.8) for $Y_n \stackrel{d}{=} \phi^n y + \sum_{j=1}^n \phi^j \eta_j$ and setting $K_0 = |\mu| + M_c$, $V_{n+1} = \sum_{j=0}^n \phi^{2j}$, choose $t = (s - \phi^n y - K_0)/(\tau^2 V_{n+1}) > 0$. Uniformly for $n \geq 1$ this yields

$$P^n q_s^{\text{ext}}(y) \leq C(1 + h)e^{-(s - \phi^n y - K_0)^2/(2\tau^2 V_{n+1})}.$$

For $b = \phi v$, the function $j(t) = \alpha(\ell - bt)^2/(1 - \phi^2 t^2)$ decreases on $[0, \phi]$: its derivative has the sign of $\phi^2 \ell t - b < 0$. Consequently $j(\phi^n) \geq j(\phi) = J_0$. Summing with weights ϕ^n bounds $R_\theta q_s - q_s$ by $C(1 + h)e^{-J_0 h^2/(2\tau^2) + Ch}$. The boundary term in the derivative equation, after applying R_θ , is at most $B\psi(dh)\|z'\|_\infty/(1 - \phi)$. Thus, with $k_0 = \min(J_0, 2d^2) > D_0^2$,

$$\|z' - q_s\|_\infty \leq C(1 + h^2)e^{-k_0 h^2/(2\tau^2) + Ch}.$$

On J_h , $p_s \geq c_1 e^{-D_0^2 h^2/(2\tau^2) - Ch}$ and $q_s \geq c_1 h e^{-D_0^2 h^2/(2\tau^2) - Ch}$. Finally $z(y) = p_s(y) - p_s(0) + \int_0^y (z' - q_s)$ and $p_s(0) \leq C e^{-\ell^2 h^2/(2\tau^2) + Ch}$. The strict gaps in (D.9) prove (D.10) uniformly, including mixtures depending on h . $\square$

D.2.3 Amplitudes bounded away from zero

Lemma D.3 (Response and paths). *For large h , uniformly over $0 < a < 1$,*

$$G(x) \geq 0 \quad (0 \leq x \leq h), \quad G(x) \geq \frac{1}{2}\psi(h - \phi x) \quad (\phi h \leq x \leq h), \quad (\text{D.12})$$

$$\int_h^U f_+ G \leq C(1 + h^3)e^{-(1-\phi)h^2/\tau^2 + Ch}. \quad (\text{D.13})$$

For every fixed $n \geq 3$, put $V_n = \sum_{j=0}^{n-1} \phi^{2j}$, $b_n = \sum_{j=0}^{n-1} \phi^j$ and $M_n = \mu b_n$. There is $c_n > 0$ with

$$f_h(h) \geq c_n e^{-(h-M_n)^2/(2\tau^2 V_n)}. \quad (\text{D.14})$$

Also, for fixed $0 < p \leq q < 1$ and $\zeta > 0$, uniformly when $U \geq h$, $ph \leq x \leq qh$, and $L \geq \phi^2 x + \zeta h$,

$$f_-(x) \geq c_n e^{-(x-M_n)^2/(2\tau^2 V_n)} \quad (\text{D.15})$$

for sufficiently large h . The cutoff may depend on n, p, q, ζ .

Proof. On $[0, b]$, $b = \phi h$, the derivative argument of Lemma D.2 applies to C_h . Put $p = p_h$, $q = p'$, $g = \Gamma'$, $\gamma = \Gamma(b)$ and $R = (I - \phi A)^{-1}$, where A is the nonreset part of C_h . Differentiation and integration over $[0, b]$ give

$$g = R(q - \gamma p), \quad \gamma = \frac{\int_0^b Rq}{1 + \int_0^b Rp} > 0.$$

For $u = (1 - \phi)h$, let $H_u = (u - \mu)/\tau^2$. Eventually $\|q\| = H_u \psi(u)$, and

$$\gamma \leq \frac{bH_u \psi(u)}{1 - \phi} = O(h^2 \psi(u)) < H_u.$$

As $q \geq H_u p$ and R is positive, $g > 0$. The Poisson equation then gives $\Gamma(y) \geq p_h(y) - m$ for $y \leq b$ and $\Gamma(y) \leq p_h(y) + \gamma$ for every $y \leq 2\phi h$. By (D.5), $m \leq Ce^{-\alpha h^2/(2\tau^2) + Ch}$, whereas $\psi(h - \phi x) \geq c_1 e^{-\alpha^2 h^2/(2\tau^2) - Ch}$ for $\phi h \leq x \leq h$. Since $\alpha > \alpha^2$, this proves (D.12). For (D.13), bound the integral by

$$\int_h^{h/\phi} f_+(x) \psi(h - \phi x) dx + \gamma \Pr(s_+ \geq h).$$

The quadratic $\alpha(x - B_0)^2 + (h - \phi x - \mu)^2$ has its minimum over $x \geq h$ at h for large h , with value $2(1 - \phi)h^2 + O(h)$. The second term has the same exponent because $\alpha + (1 - \phi)^2 = 2(1 - \phi)$. Equation (D.5) proves the bound.

For the path bounds, at least half of either incoming carry law lies in $[0, 2M_c]$. Condition on y there and on an n -inflow path ending at x in the low phase. Put $w_j = \phi^{n-j}$ and

$D_x = x - \phi^{n-1}y - M_n$. Center the inflows at $\xi_j = \mu + w_j D_x / V_n$; integrate over

$$\eta_j = \xi_j + t_j \quad (j < n), \quad |t_j| \leq \tau/4, \quad \eta_n = \xi_n - \sum_{j < n} w_j t_j.$$

The endpoint is x with Jacobian one, and all inflows are positive uniformly for large h . For $j < n$ the intermediate raw stock is

$$S_j = \phi^{j-1}y + \mu b_j + \frac{\phi^{n-j}V_j}{V_n}D_x + \sum_{l \leq j} \phi^{j-l}t_l \leq \phi^{n-j}x + C_n.$$

Low-phase intermediate dates have even $n - j \geq 2$, so $S_j < L$; high-phase dates have $S_j \leq \phi q h + C_n < h \leq U$. Every intermediate decision avoids reset; the final decision is unrestricted. The perturbation vector is orthogonal to (w_j) , giving

$$\sum_{j=1}^n (\eta_j - \mu)^2 = D_x^2 / V_n + \sum_{j < n} t_j^2 + \left(\sum_{j < n} w_j t_j \right)^2 \leq (x - M_n)^2 / V_n + C_n.$$

Integrating over this fixed box and half the incoming mass proves (D.15) with a positive c_n depending only on n, μ, τ, ϕ . For $x = h$ at the fixed threshold, every intermediate stock is at most $\phi h + C_n < h$, proving (D.14). $\square$

Use the single boundary $r_0 = (\phi^3 + \phi^2)/2$, so $\phi^3 < r_0 < \phi^2$. The following bounds hold uniformly for $a \geq a_s$.

Case $r = L/h \geq r_0$. Here $r \leq 1 - a_s$. Put

$$m_0 = \min(\alpha r_0, r_0 - \phi^3) > 0, \quad \delta = \min\{a_s/8, m_0/(8\phi^2)\}, \quad q = 1 - a_s/2.$$

The interval $I = [(\max(r, \phi) + \delta)h, (\max(r, \phi) + 2\delta)h]$ lies in $[L, h] \cap [\phi h, qh]$ and satisfies $L - \phi^2 x \geq (m_0/2)h$. Let $g_0 = 2(1 - \phi) - \{\alpha + (q - \phi)^2\} > 0$ and choose a fixed $n \geq 3$ with $V_n^{-1} - \alpha < g_0/2$. For $\rho = x/h \in [\phi, q]$,

$$\rho^2 / V_n + (1 - \phi\rho)^2 = \alpha + (\rho - \phi)^2 + (V_n^{-1} - \alpha)\rho^2 \leq 2(1 - \phi) - g_0/2.$$

Equations (D.12) and (D.15), on a strip of width δh , imply

$$\int_L^h f_- G \geq \frac{1}{2} \int_I f_-(x) \psi(h - \phi x) dx \geq c_1 h e^{-(1-\phi)h^2/\tau^2 + g_0 h^2/(4\tau^2) - Ch}. \quad (\text{D.16})$$

The discarded part is nonnegative by (D.12).

Case $0 < r = L/h \leq r_0$. Let $v_2 = 1 + \phi^2$ and

$$g_2(t) = (\tau \sqrt{2\pi v_2} Z^2)^{-1} e^{-(t - (1+\phi)\mu)^2 / (2\tau^2 v_2)}.$$

The raw equations and untruncated two-normal convolution give

$$\begin{aligned} f_+(h) &\leq \psi(h) + \int_0^L f_-(z) \psi(h - \phi z) dz, \\ f_-(h) &\leq \psi(h) + g_2(h) + \int_0^L f_-(z) g_2(h - \phi^2 z) dz. \end{aligned}$$

The terms $\psi(h)$ and $g_2(h)$ account for reset at the two possible decisions. The convolution bounds extend only nonnegative integrals. Using (D.5) and

$$\alpha z^2 + (h - \phi z)^2 = \alpha h^2 + (z - \phi h)^2, \quad \alpha z^2 + (h - \phi^2 z)^2/v_2 = \alpha h^2 + (z - \phi^2 h)^2/v_2,$$

we get, uniformly for $L \leq r_0 h$,

$$F(h) \leq C(1 + h)e^{-(\alpha + k_2)h^2/(2\tau^2) + Ch}, \quad k_2 = (\phi^2 - r_0)^2/v_2 > 0. \quad (\text{D.17})$$

Indeed $(\phi - r_0)^2 > k_2$, while the innovation-only terms have gaps ϕ^2 and ϕ^4/v_2 , also larger than k_2 . The bound is uniform as L/h tends to zero.