A Theory of Lumpy Innovation and Aggregate Dynamics*

Lixing Wang† Donghai Zhang‡

May 2, 2025

Abstract

We propose a new theory of lumpy innovation by developing a business cycle model with endogenous innovation, heterogeneous innovating firms, and fixed adjustment costs. Through this framework, we uncover a novel mechanism underpinning the "Great Twins" hypothesis: the Great Moderation–marked by diminished economic volatility up to 2008–causally intensified the severity of the subsequent Great Recession. Moreover, the same mechanism leads to the underutilization of innovative ideas. As a result, reduced volatility can hinder average U.S. economic growth. The paper also provides both micro- and macro-level evidence to substantiate the proposed theory of lumpy innovation.

Keywords: Heterogeneous Innovators, Lumpy R&D, Distribution, Great Twins, Medium-run Business Cycles

JEL Classification: E32, O31, O33

*We thank Diego Comin, John Fernald, Pablo A. Guerrón-Quintana, Franck Portier, Yongseok Shin, David Lagakos, Pablo Ottonello, Michael Song, 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'.

†University of Bonn, lixing.wang@uni-bonn.de

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

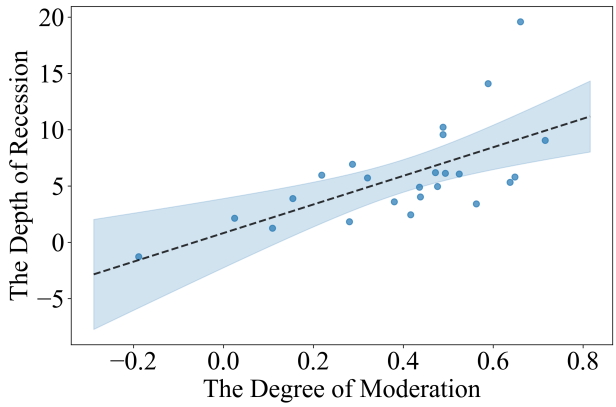
1 Introduction

Research and Development (R&D) create new knowledge and technology that promote economic growth. Business cycles affect the profitability of new technology and, therefore, influence innovators' innovation decisions. This insight, coupled with the observed pro-cyclicality of *aggregate* R&D expenditure, has motivated a growing literature (Comin and Gertler, 2006) that emphasizes the effects of business cycles on medium-run economic growth via their impact on innovation (or R&D expenditures). This theoretical literature has provided valuable insights, focusing on homogeneous innovators and emphasizing changes in innovators' decisions along the intensive margin.

We consider an important yet unexplored *micro-level* feature: changes in R&D expenditures are *lumpy* rather than smooth. This lumpiness is an indication of the relevance of innovators' innovation decisions along the extensive margin: to set up or not a new R&D project. Therefore, business cycle fluctuations can have medium-run effects via their impact on innovation decisions along the extensive margin. More importantly, the strength of these effects depends on the number of innovators at the margin of initiating new R&D projects and the amount/quality of innovative ideas they possess. In other words, the severity of a medium-run recession depends on the *distribution* (mass) of innovators positioned at the threshold of launching a new R&D project. This distribution can be affected by macroeconomic policies, and therefore, policies might have unintended consequences affecting the left-tail risk. Motivated by these insights, we ask: What are the (policy) implications of lumpy innovation for medium-run business cycles?

Theoretically, we develop a business cycle model with *heterogeneous* innovators who face fixed adjustment costs when setting up new R&D projects, giving rise to lumpy innovation. The model offers a novel mechanism (Lemma 1): the asymmetric effects of shocks on the distribution of innovators at the adjustment threshold. As a result, macroeconomic *volatility* alters the *shape* of this distribution (Proposition 2). This mechanism gives rise to a novel insight: the Great Twins hypothesis. Specifically, the Great Moderation–characterized by reduced economic fluctuations until 2008–causally contributed to the observed severity of the subsequent 2008 Great Recession. The Great Twins hypothesis highlights a new tail-risk and volatility tradeoff faced by central banks: while effective monetary policy reduces unconditional volatility, it amplifies the GDP decline conditional on a recession. Additionally, the mechanism implies that moderation periods are associated with an over-accumulation of ideas. When ideas are subject to depreciation, periods

Figure 1: The Great Twins in the Data



Notes: This figure plots the depth of the Great Recession for each country–defined as the gap between actual GDP and the GDP predicted by the pre-recession trend–against the degree of moderation, measured as the ratio of the standard deviation of GDP growth over the 15 years preceding the recession to that over the full sample. The analysis is based on cross-country data from the OECD, sourced from the World Bank. Further details are provided in Section 7.

of moderation can lead to lower macroeconomic growth.

Empirically, we present macro-level evidence supporting the Great Twins hypothesis (Figure 1), micro-level evidence consistent with the mechanism behind the Great Twins, and additional evidence validating the relevance of innovation along the extensive margin.

The insights and findings presented in this paper are relevant for both academics and practitioners. A reduction in aggregate fluctuations is conventionally viewed as desirable. However, this paper introduces a novel mechanism through which reduced fluctuations amplify tail risk and lead to underinvestment in innovation. It is important to note that fluctuations and volatility in this paper refer to *realized* sample volatility, which is distinct from uncertainty.[1]

In more detail, we develop a business cycle model with endogenous innovation following the framework of Romer (1990) and Comin and Gertler (2006). Our model features heterogeneous innovators facing a fixed setup cost for initiating new R&D projects. The novelties of our model lie in the innovators' problems, which unfold in three phases. The first phase is termed the innovation phase, during which each innovator receives/improves raw ideas. Innovators are heterogeneous in terms of the amount of raw ideas they have.

[1]Therefore, our mechanism differs from standard uncertainty-based channels, such as wait-and-see behavior or risk-aversion-driven responses.

In the model, the amount of ideas is isomorphic to the quality of an idea. The function of an R&D project is to process the stock of raw ideas into R&D ideas. We distinguish between raw ideas and R&D ideas, where the latter refers to ideas actively pursued within an on-going R&D project. To create a new R&D project (to innovate), the innovating firm incurs a fixed cost. The key decision innovators make during the innovation phase is whether to set up a new R&D project, i.e., the extensive margin innovation decision. To convert R&D ideas into new technologies, the source of growth, innovators need to hire scientists. Scientists' wage bills constitute the variable cost. This describes the second phase, called the adoption phase. In the last phase, the monetizing phase, innovators claim the intellectual property right (patent) of the new technologies and sell the patent to an intermediate input producer.

The model incorporates standard business cycle features that permit the analysis of the interaction between business cycle fluctuations and innovation decisions. Business cycle fluctuations affect the profits of new technologies and, therefore, alter innovators' decisions, particularly along the extensive margin. The model predicts that R&D spike adjustments are pro-cyclical, consistent with our empirical observations. The model is calibrated to match key moments in the U.S. economy and is shown to match untargeted moments, including the motivating facts described above, and the distribution of innovators. Furthermore, we fit the model to the evolution of real GDP in order to estimate business cycle shocks and underlying distributions, which serve as inputs for the counterfactual exercises. Next, we turn to the main implications of the model.

The model offers a novel *qualitative* insight: the Great Moderation that pre-dated the Great Recession has contributed to the severity of the latter — a phenomenon labeled as the *Great Twins*. Specifically, we show that a decrease in the sample volatility of GDP in the years preceding the Great Recession, either due to Good Luck (drawing smaller shocks) or Good Policy (U.S. Federal Reserve adopted a more effective monetary policy to reduce economic fluctuations), *causally* increases the effects of shocks that triggered the Great Recession on productivity and real GDP. The Great Twins hypothesis has an important implication for policymakers: a well-intentioned policy that reduces macroeconomic volatility may inadvertently increase left-tail risk. Relatedly, it suggests that macroprudential policies can backfire by amplifying vulnerabilities that only materialize during rare but severe downturns.

The key mechanism underpinning the Great Twins hypothesis arises from the asymmetric effects of shocks on the distribution of raw ideas (Lemma 1). A positive shock

leads to more innovation along the extensive margin–that is, greater utilization of raw ideas. As a result, fewer raw ideas remain available in the subsequent period: positive shocks compress the distribution. Conversely, negative shocks lead to an accumulation of ideas, thereby stretching the distribution. We show analytically that when innovators are more likely to draw a small idea than a large one (or, equivalently, a mediocre idea rather than a great one), a positive shock has a greater absolute effect on the distribution than a negative shock of the same magnitude. This asymmetry implies that reduced macroeconomic volatility stretches the distribution of total ideas across innovators, making it *more fat-tailed* (Proposition 2). Consequently, innovators are more likely to adjust along the extensive margin, increasing the likelihood of extreme adverse outcomes associated with declining innovation. To the best of our knowledge, this mechanism has not been documented before.

Due to the same mechanism, an economy experiencing greater moderation shifts the distribution of innovation ideas to the right–that is, it overaccumulates (or underutilizes) raw ideas. Under the assumption that unused raw ideas are subject to depreciation, this moderation-induced overaccumulation (or underinvestment) is associated with lower average GDP growth.

We present empirical evidence that validates the model. First, we provide macro- and micro-level evidence supporting the Great Twins hypothesis and validating its underlying mechanism. Using cross-country variation, we show that countries which experienced more economic moderation in the decade preceding the 2008 crisis also suffered deeper recessions during the Great Recession (Figure 1). This finding is not unique to the 2008 episode: extending the empirical analysis to all post-WWII recessions yields a similar pattern.

Second, the model suggests that the shape of the distribution before the Great Recession shocks is crucial for the observed depth of the Great Recession. We confirm this observation using micro-level R&D expenditure data.

Third, using micro-level R&D expenditure data, we document that the distribution of changes at the firm level is heavily populated by zeros/small values and large values–a lumpy pattern. Moreover, the fraction of firms making large (or spike) adjustments in R&D expenditure–the spike rate–is procyclical. This evidence supports the existence of innovation decisions along the extensive margin and highlights their cyclical nature.

Fourth, by merging firm-level R&D data with patent data, we show that a firm's extensive adjustment in R&D expenditure is positively correlated with future increases in

its patenting activity, confirming the material impact of extensive-margin innovation decisions.

Related Literature The theoretical part of our paper contributes to the medium-term business cycle literature, which introduce business cycles to Romer (1990)'s endogenous growth model: see, e.g., Comin and Gertler (2006), Comin et al. (2009), Kung and Schmid (2015), and Gornemann et al. (2021).[2] Methodologically, we contribute this line of research by introducing two important features: (i) heterogeneous innovators whose R&D investment decisions are subject to (ii) a fixed cost of launching new R&D projects. These new features allow us to understand how business cycle *fluctuations* affect the distribution of innovators' propensity to innovate and, importantly, how these distribution changes influence the aggregate dynamics.[3] Related to this insight but through a different mechanism, Bilbiie et al. (2012) and Bilbiie and Melitz (2020) study how firm entry and exit affect the propagation of business cycle shocks.

There is a long tradition of analyzing the aggregate implications of firms' lumpy capital investment behavior with fixed adjustment costs; see, for example, Caballero et al. (1995), Caballero and Engel (1999), Khan and Thomas (2008), Bachmann et al. (2013), Miao and Wang (2014) and more recently, Koby and Wolf (2020), Winberry (2021), Baley and Blanco (2021), Gnewuch and Zhang (2025), Chiavari and Goraya (2024) and Lippi and Oskolkov (2023). To the best of our knowledge, the new mechanism in our model has not been previously documented in the literature. Unlike the typical lumpy capital investment (Ss) model, our lumpy innovation exhibits a one-S feature, which is crucial to the mechanism behind the Great Twins: the asymmetric effects of shocks on the distribution.[4] Secondly, since the incentive for innovation primarily comes from innovation profits, our model

[2]This approach is different from Matsuyama (1999), who connects business cycle fluctuations with growth through alternating phases of high investment and high innovation. Differing from Matsuyama (1999), we present a novel model that connects business cycle fluctuations with innovators' R&D decisions along the extensive margin.

[3]By introducing firm heterogeneity into a model with endogenous growth, our framework complements endogenous growth models featuring Schumpeterian innovation, see, e.g., Aghion and Howitt (1992), Klette and Kortum (2004), Acemoglu et al. (2018), and Akcigit and Kerr (2018). Complementing these papers, our framework enriches the dynamics of R&D by modeling the accumulation of innovative ideas and the role of fixed adjustment costs, thereby extending the analysis beyond static determinants such as firm size or the number of product lines. Our framework highlights heterogeneity in the propensity to innovate and generates implications beyond variables in the steady state. Bloom (2007), Brown and Petersen (2011), and Chen et al. (2021) investigate the implications of *convex* adjustment costs on R&D investment, we instead focus on the role of a *fixed* (non-convex) component in adjustment costs.

[4]In an Ss model, adjustments can occur in either direction, so aggregate shocks do not generate the kind of stretching or compressing of the distribution emphasized above. This partly explains why the mechanism we emphasize has not been uncovered before.

does not encounter the issue of irrelevance of fixed costs due to the semi-elasticity of investment with respect to the real interest rate approaching infinity, as seen in the capital investment literature (see, e.g., Caballero and Engel 1999, Khan and Thomas 2008, Koby and Wolf 2020, Winberry 2021). In terms of insight, our paper contributes to this class of models by showing that realized volatility alters aggregate dynamics by affecting the shape of the distribution.

The empirical section of our study is related to the literature examining the cyclicality of R&D. Fatas (2000), Barlevy (2007), Broda and Weinstein (2010), Ouyang (2011), and Fabrizio and Tsolmon (2014) document that aggregate R&D spending is pro-cyclical, with this pro-cyclicality being observed across different contexts and for various sources of R&D data and cyclical measures. The pro-cyclicality of R&D is also evidenced by patterns in patenting (Griliches 1990) and technology diffusion (Anzoategui et al. 2019), with Comin and Gertler (2006) further revealing its presence at both short- and medium-term frequencies. Our paper contributes to this body of literature by highlighting the importance of R&D decisions along the extensive margin, their cyclical nature, and their significant impact on patenting outcomes.

Last but not least, this paper contributes to the growing literature on the Great Recession and R&D investments. Benigno and Fornaro (2018) and Anzoategui et al. (2019) attribute the post-Great Recession slowdown in productivity growth to a persistent decline in technology creation and adoption, driven by the zero lower bound on nominal interest rates. Queralto (2020) and Guerron-Quintana and Jinnai (2019) introduce financial frictions to the medium-run business cycle model by Comin and Gertler (2006) and explore how these frictions hinder innovation, leading to a prolonged economic slowdown. Additionally, Bianchi et al. (2019) and Schmitz (2021) incorporate financial frictions into a vertical innovation framework to explain the long-term impacts of the 2008 recessions. We offer a new explanation, the Great Twins. Our framework can explain the joint and sequential occurrences of the Great Moderation and the Great Recession. Importantly, the nonlinear mechanism—the causal effects of reduced volatility (a sequence of small shocks) on increased left-tail risk—in our model is an *endogenous* feature. Moreover, different from the existing literature, we emphasize R&D investment along the extensive margin. Overall, we view the main contribution of our paper as providing a new qualitative insight.

The rest of the paper is organized as follows. Section 2 presents motivating evidence on lumpy innovation. Section 3 illustrates the mechanism in a partial equilibrium setting,

which is then embedded in general equilibrium in Section 4. Section 5 describes the model calibration. Section 6 presents the main results. Section 7 provides both micro- and macro-level evidence to support the proposed theory of lumpy innovation.

2 Motivating Evidence

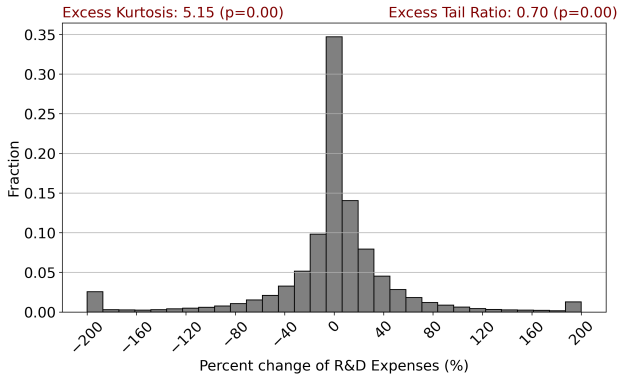
This section presents novel empirical evidence that motivates the model introduced in the subsequent sections. Specifically, we document that the *changes* in R&D investments of innovating firms exhibit a lumpy pattern characterized by a substantial proportion of spike adjustments and inaction.

Data Description Our empirical analysis leverages annual firm-level panel data from S&P Capital IQ's Compustat database, which includes consolidated financial data from both domestic and international subsidiaries for the period from 1968 to 2020. We specifically focus on U.S. firms, excluding those in the financial and utility sectors due to their high regulation and capital requirements. Some firms never engage in R&D activities, and our analysis concentrates on innovating firms. In the empirical analysis, a firm is classified as an innovating firm if it reports positive R&D expenses at least once in our sample period. For cross-time analysis, we retain firms with a minimum of five observations. We also exclude acquisitions exceeding 5% of total assets.[5] Our sample consists of 7,214 innovating firms, representing 20.4% of the companies in the Compustat universe. While our sample covers a similar set of industries as the full Compustat population, it places greater weight on manufacturing firms, which are more R&D-intensive, as detailed in Table B.1 in the Online Appendix. The panel is unbalanced, with each firm averaging 18.9 observations.

The primary variable in our study is R&D Expense, encompassing all costs, including salaries and departmental expenses, related to the development of new products or services during the year. R&D expense is reported as a supplementary item on the Income Statement. Besides salaries–the main source of R&D expense–and other costs associated with the development, design, and testing of new products or improvement of existing products, which take effect immediately, it includes amortized software costs and other capitalized R&D assets. R&D expenses are deflated using the GDP deflator to obtain their real counterparts. Henceforth, R&D expenses refer to their real values.

[5]Mergers and acquisitions may incorporate the R&D departments of acquired companies, causing abrupt changes in R&D expenditures.

Figure 2: Lumpy Innovation



Notes: This figure plots firm-level distribution of percent change of R&D expenses. To handle cases with zero initial R&D expenses, we use the formula: percent change $= \frac{\text{R\&D}_t - \text{R\&D}_{t-1}}{0.5 \times (\text{R\&D}_t + \text{R\&D}_{t-1})}$. Excess kurtosis measures deviation from a normal distribution's fourth moment. The excess tail ratio compares the 80%-to-95% quantile ratio to that of a normal distribution. P-values are obtained via bootstrapping.

Formalizing the Concept of Lumpy Innovation To formalize the concept of lumpy innovation, Figure A.1 in the Online Appendix presents an example of how R&D expenses evolve over time for an actual firm in the data. Most of the time, *changes* in R&D expenses remain small. However, occasionally, this firm makes a spike adjustment in its R&D. We refer to these instances as evidence of lumpy innovation.

An innovation, or starting a new R&D project, requires a significant change in R&D expenditure both today and in the future due to, for example, the hiring of new scientists and their associated wage bills. Formally, lumpy innovation postulates that the distribution of *changes* in R&D is heavily populated by zeros/small values and large values. Is lumpy innovation a general phenomenon in the data? This question is addressed in the following section.

Motivating Fact: Lumpy Innovation Figure 2 plots the distribution of the the percent change of R&D expenses. It is apparent that firms' R&D investment decisions satisfy the properties of lumpy innovation: R&D investment is characterized by no or small changes and large spikes. The excess kurtosis of this distribution is 5.15 which implies a fat-tail pattern of the distribution and strongly rejects the hypothesis of a normal distribution with a p-value of 0.00. Furthermore, 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%, significantly exceeding the null value of zero and confirming the presence of a fat-tail pattern.

Robustness Checks In Online Appendix B.2, we provide a detailed discussion of additional robustness checks for this fact. First, we show that the lumpy pattern observed in Figure 2 is not simply a reflection of volatile changes in firms' scale[6] and we also confirm that the lumpy pattern holds for patent data. Second, we establish that this lumpiness is consistent across firms of different R&D expenditure sizes, addressing concerns that our findings might be disproportionately influenced by smaller R&D firms. Third, we demonstrate that lumpy changes in R&D are not driven by financial frictions: the observed pattern remains qualitatively and quantitatively similar across firms with both high and low leverage. Fourth, we argue that strategic accounting for tax purposes does not drive the lumpy pattern—U.S. R&D tax credit rules, in fact, discourage it. Lastly, the pattern is not due to entry or exit: the excess kurtosis is 4.94 ($p = 0.00$), and the excess tail ratio is 0.51 ($p = 0.00$) when entry and exit observations are excluded.

In summary, the lumpiness of innovation that we emphasize stresses the significant presence of two key regions in the distribution of changes in R&D expenses: the inaction region—characterized by zero or small changes in R&D expenses, and the spike region—representing spike adjustments. In the remainder of the paper, we provide a model that interprets the spike adjustment in R&D as evidence suggesting that the innovator is setting up a new innovation project—R&D investment along the extensive margin.

3 Heterogeneous Innovators: Key Mechanisms

We emphasize *heterogeneous* innovating firms that must pay a fixed cost to set up a new R&D project. Incorporating heterogeneous agents increases the model's complexity by introducing the firm-level distribution of unadopted technologies–referred to as the idea stock distribution–as a state variable, which is a central focus of this paper.

The innovation block in our model builds upon Romer (1990). There is a continuum of innovators with a total measure of unity. Each innovator *creates* new varieties of intermediate inputs and profits by selling the associated Intellectual Property Rights. Figure A.2 provides an overview of the innovation block, which consists of three phases within each period. In Phase I, innovators receive new raw ideas and decide whether to initiate a new R&D project, drawing on their accumulated stock of raw ideas. In Phase II, ideas embedded in existing R&D projects—henceforth referred to as R&D ideas—undergo a process of

[6]Specifically, by documenting the lumpiness of the R&D-to-sales ratio, we show that lumpy innovation is not driven by jumps in sales. Moreover, we demonstrate that the lumpy pattern persists even among firms with relatively stable sales.

random adoption. This phase involves hiring skilled labor for idea adoption, with R&D expenses primarily reflecting the labor costs of these workers. Finally, in Phase III, successfully adopted R&D ideas materialize into new varieties of intermediate inputs. The innovator then monetizes these developments by selling the Intellectual Property Rights of the newly created varieties to the intermediate input producer, as described in Section 4.2.

It is worth emphasizing that ideas can be adopted and become a new variety *only if* they are embedded in an R&D project. Therefore, we distinguish between raw ideas and R&D ideas. In fact, the central focus of this paper is to examine innovators' extensive margin decision in Phase I: whether to set up a new R&D project that transforms raw ideas into R&D ideas.

The main insight of the paper is that the distribution of idea stocks, which depends on *realized* fluctuations in the economy, plays a key role in shaping aggregate dynamics. To emphasize the qualitative nature of the key mechanisms, we begin by deriving propositions in a partial equilibrium setting that analyze innovators' optimization in Phase I. Section 4 will complete the description of innovators' optimization in Phases II and III and incorporate these features into a general equilibrium model with business cycle shocks.

3.1 Innovators' Decisions in Phase I

Consider an economy with a continuum of heterogeneous innovators, each accumulating raw ideas over time. The stock (or quality) of raw ideas, denoted by $s_{i,t}$, evolves dynamically. Innovators face a control problem in which they can choose whether to set up a new project that transforms their accumulated raw ideas into new technologies, subject to a fixed setup cost c.

When an innovator does not adjust, the stock of ideas evolves according to:

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

where $0 \leq 1 - \phi \leq 1$ represents the rate at which ideas become obsolete, and $\eta_{i,t}$ captures the inflow of new ideas. To analytically illustrate the key mechanism, we assume in this section that $\eta_{i,t}$ is exogenous and drawn from an independent and identically distributed process. We further assume that $\eta_{i,t}$ is non-negative, capturing the additive nature of knowledge accumulation, and that its probability density function is continuous and strictly decreasing. This reflects the idea that small improvements are more common

than large breakthroughs. We relax these assumptions in Section 4. 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.

Each unit of a raw idea has a chance to be incorporated into an innovation project, which can subsequently be converted into a profitable new variety. Therefore, ideas are valuable. Let $\tilde{J}_t$ denote the unit value of raw ideas:

$$\tilde{J}_t = \tilde{J}(\varrho_t). \tag{3.2}$$

In this partial equilibrium setting, we assume that $\tilde{J}_t$ is a function of exogenous shocks ϱ_t. This relationship will become endogenous in the general equilibrium framework discussed in Section 4.

An innovator's problem is to choose between adjusting and not adjusting, based on her current stock of ideas. Denoting $S_t(s_{i,t})$ as the value of innovator i with a stock of ideas $s_{i,t}$ in period t, this problem can be written as:

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

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.

In the absence of convex adjustment costs, innovators incorporate all their accumulated ideas at once when they decide to set up a new project. The value of an innovator i who decides to adjust is:

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

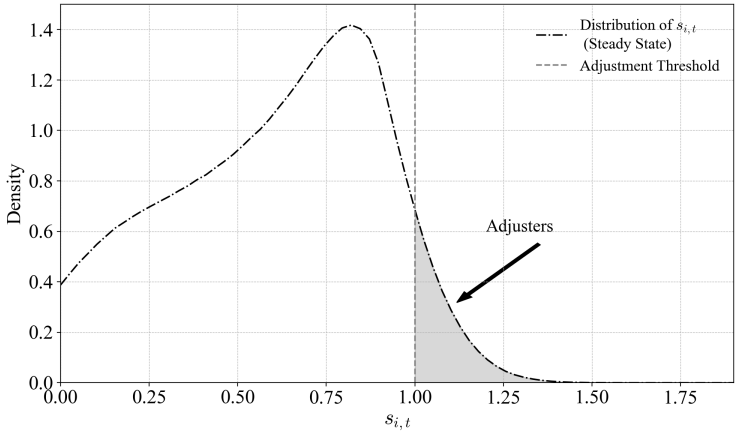
where the first term represents the adjustment cost, $s_{i,t}\tilde{J}_t$ captures the payoff from the current stock of ideas, and $E_t\left[\Lambda_{t,t+1}S_{t+1}(\eta_{i,t+1})\right]$ is the expected continuation value of the innovator if they reset their stock of ideas to zero.

The continuation value of an innovator i who decides not to adjust in the current period is:

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

Solving the innovators' problem, as described in Equations (3.3)–(3.5), gives rise to the following proposition.

Figure 3: R&D Decisions Along the Extensive Margin



Proposition 1. *If $c > 0$, there exists a threshold $\bar{s}_t > 0$ such that innovators choose to adjust if and only if their stock of ideas exceeds $\bar{s}_t$.*

Proof. See Online Appendix D. □

Figure 3 visualizes Proposition 1. It plots the distribution of $s_{i,t}$ together with the threshold threshold $\bar{s}_t$. The shaded area indicates the aggregate number of adjusters—innovators who decide to set up new projects—in the economy.

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

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

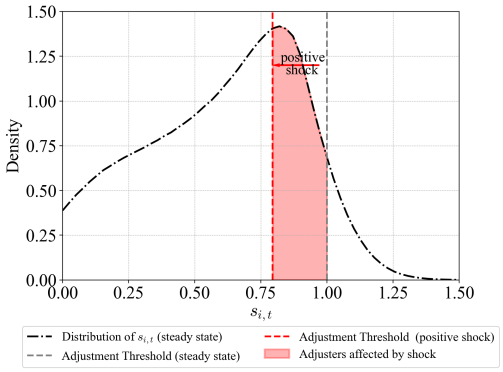
where $I\{s_{i,t} < \bar{s}_{i,t}\}$ is an indicator variable that equals 1 if $s_{i,t} < \bar{s}_t$ and 0 otherwise.

It is evident that the overall level of innovation in the economy depends on the distribution of raw ideas and the threshold $\bar{s}_{i,t}$. Business cycle shocks affect both factors, thereby influencing aggregate innovation. Next, we illustrate how business cycle shocks and fluctuations impact the distribution of raw, unused ideas—the key mechanism driving the main results of the paper.

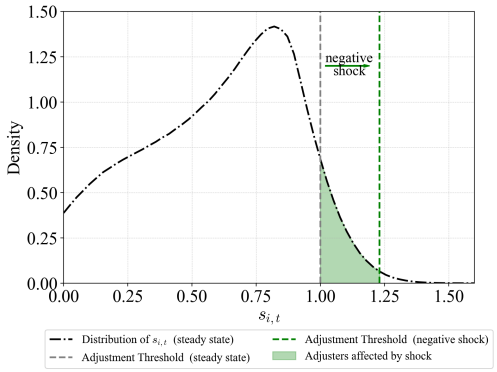
3.2 Mechanisms: Fluctuations and Distributions

Consider a positive shock that increases the value of raw ideas, $\tilde{J}_t$. This incentivizes innovators to accelerate the implementation of their accumulated raw ideas, leading to a leftward shift in the adjustment threshold, as shown in Figure 4a. As a result, more in-

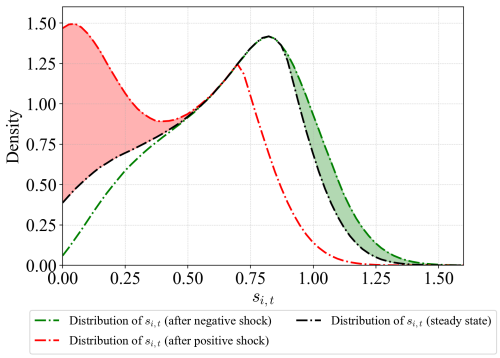
Figure 4: Illustration of Shocks and Effects on the Distribution of $s_{i,t}$



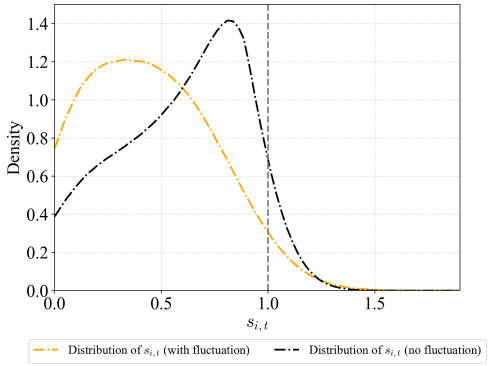
(a) Positive shocks shift $\bar{s}_{i,t}$ to the left



(b) Negative shocks shift $\bar{s}_{i,t}$ to the left



(c) Asymmetric Effects on the Distribution



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

novators with a sufficiently large stock of ideas choose to adjust, depleting their accumulated stock, while fewer continue to accumulate. This, in turn, compresses the distribution of $s_{i,t}$ in the subsequent period, as illustrated by the red dashed line in Figure 4c.

Conversely, a negative shock reduces the profitability of raw ideas, leading innovators to delay the implementation of raw ideas until the economy recovers. As a result, the adjustment threshold shifts rightward, as shown in Figure 4b. This shift discourages R&D investment during the downturn, causing more innovators to continue accumulating their stock of innovative ideas. In turn, this accumulation process stretches the distribution of $s_{i,t}$, as illustrated by the green dashed line in Figure 4c.

This paper emphasizes tail risk, defined as the amount of economic growth forgone due to a big negative shock. This tail risk is proportional to the mass of innovators near the adjustment threshold $\bar{s}^*$. Therefore, changes in the distribution of $s_{i,t}$, whether through compression or expansion, affect this risk. Under this interpretation, Lemma A.1 pro-

vides a formal representation of the compressing and stretching effects: by compressing (stretching) the distribution, a positive (negative) shock lowers (raises) the economy's underlying risk in the subsequent period.

The following lemma shows that, in the model of lumpy innovation—characterized by a one-sided adjustment threshold, with a non-negative idea inflow whose density decreases in inflow size—the compression from positive shocks (red shaded area) exceeds the stretching from negative shocks (green shaded area), which leaves a smaller mass around the adjustment threshold $\bar{s}^*$–hence smaller tail risk–given the presence of the fixed adjustment cost.

Lemma 1. *In the model with lumpy innovation, a positive shock to the adjustment threshold induces a greater absolute change in the distribution of $s_{i,t}$ than a negative shock of the same magnitude.*

Proof. See Online Appendix D. □

Lemma 1 presents the result based on a one-shot shock. The same intuition applies when considering a sequence of shocks. Proposition 2 generalizes it to study how fluctuations–sequence of shocks–affect the distribution and tail-risk. Specifically, for simplicity, the sequence of shocks ϱ_t follows a simple harmonic motion, such that the threshold $\bar{s}_t$ fluctuates around its steady-state value $\bar{s}^*$ according to $\bar{s}_t = (1 + (-1)^t \sigma)\, \bar{s}^*$. This assumption is relaxed in the quantitative general equilibrium model presented in Section 4. Proposition 2 shows that increased volatility (bigger σ) reduces the mass around the adjustment threshold, i.e., tail risk.

Proposition 2. *In the model with lumpy innovation, the mass around the adjustment threshold $\bar{s}^*$ in the stationary distribution satisfies:*

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

where f^* denotes the steady-state distribution of $s_{i,t}$ in the absence of shocks, and f_σ denotes the stationary distribution under fluctuations of magnitude $\sigma > 0$.

Proof. See Online Appendix D. □

Figure 4d illustrates Proposition 2 graphically. When the economy is subject to a sequence of shocks, the stationary distribution of $s_{i,t}$ in the fluctuating environment exhibits less mass around the extensive margin compared to the steady-state distribution.

Discussions The intuition behind the key mechanism is as follows. For the asymmetric effects of positive and negative shocks on the idea distribution to arise, the adjustment threshold, $\bar{s}$, must lie in the declining region of the steady-state distribution of ideas. Three features of the model determine the existence and strength of the asymmetry we highlight.

The first feature is the distribution of idea *inflows*. Intuitively, if innovators received infinitely large new ideas each period, the adjustment threshold $\bar{s}$ would not lie in the declining region of the steady-state distribution of ideas. In this partial equilibrium setting, we assume the idea inflow is exogenous–an assumption we will relax in the general equilibrium model. Lemma 1 demonstrates that the asymmetry arises inherently as long as there is a decreasing likelihood of acquiring additional ideas as η increases, given the presence of fixed adjustment costs. A decreasing likelihood of acquiring additional ideas as η increases implies that innovators are more likely to draw a small η than a large one—or, equivalently, more likely to draw a mediocre idea than a great one.

Second, the fixed adjustment cost c plays a critical role. In the limit, when there is no fixed adjustment cost (i.e., $c = 0$), the extensive margin decision vanishes, and the mechanism we emphasize disappears. To further illustrate the importance of lumpiness in generating the model's mechanism, Online Appendix C.4 defines asymmetry as the difference between the red and green areas in Figure 4c. We analytically show that, for two specific idea inflow processes, the degree of asymmetry increases with higher fixed adjustment costs (i.e., greater lumpiness). Our numerical analysis confirms this pattern across a range of idea inflow processes and model parameters.

Third, the inherently one-sided adjustment feature of this theory of lumpy innovation is another key determinant. To illustrate this, Online Appendix C.6 examines whether the same mechanism holds in a lumpy capital investment model with two-sided adjustments (i.e., an *Ss* model). Briefly, since capital adjustments can occur in either direction, aggregate shocks do not lead to the stretching or compressing of the distribution emphasized above. The intuition is that, while the right-hand threshold generates the same asymmetry as in our model, the left-hand threshold introduces an asymmetry of the opposite sign, thereby offsetting the effect.

4 The General Equilibrium Model

This section embeds the key mechanisms and results described above in a partial equilibrium framework into a general equilibrium (GE) setting, which is broadly categorized into three main blocks: the business cycle, production, and innovation blocks. The GE model microfounds the connection between business cycle shocks and the value of ideas $J(\varrho_t)$. The GE model also serves as a laboratory to evaluate the quantitative importance of our channels and to conduct counterfactual policy analysis.

Section 4.1 completes the description of the innovation block. Section 4.2 covers the production block. Section 4.3 concludes the model description with the business cycle block. Section 4.4 presents the corresponding market clearing conditions that define the general equilibrium.

4.1 The Innovation Block

Phase I: Innovation Decisions Along the Extensive Margin Innovators' decisions in Phase I were described in Section 3. We begin by outlining two modifications in the GE setting.

First, the fixed setup cost is now assumed to be proportional to the number of varieties A_t, so that the fixed cost is given by cA_t. This assumption reflects that the cost of setting up an R&D project increases with technological complexity. More importantly, since the economy grows at the rate of A_t, this assumption ensures that the adjustment cost friction does not vanish over time.

Second, we relax the assumption that idea inflow $\eta_{i,t}$ is exogenous. Instead, in the GE model, an innovator i can increase idea inflows in the subsequent period by hiring skilled labor to develop ideas in the current period. Specifically, idea inflow follows:

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

where $N_{i,t}$ denotes skilled labor, and $\varepsilon_{i,t}^{\eta}$ is drawn from an independent and identically distributed distribution.[7]

These two modifications affect $S_t^a(s_{i,t})$ and $S_t^n(s_{i,t})$, as defined previously by Equations (3.4) and (3.5). In the GE model, the value of an innovator i who decides to set up a new

[7] In an earlier version of this paper, the inflow of ideas was assumed to be exogenous. All results presented in the current paper remain qualitatively the same and quantitatively similar under this assumption.

project is:

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

where W_t^{RD} denotes the wage of skilled labor. The continuation value of an innovator i who decides to wait and not set up a new project in the current period is:

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

The remaining equations in Section 3.1 remain unchanged.

Phase II: The Adoption of Ideas Embedded in R&D Projects Once a new R&D project is established, it becomes a pending project awaiting construction. In each period, a proportion $\tilde{\lambda} < 1$ of pending projects are realized and commence construction. This assumption captures the reality that R&D projects do not materialize instantaneously; it also reflects potential convex adjustment costs associated with R&D investment (Bloom 2007, Brown and Petersen 2011, Chen et al. 2021), including delays due to securing funding, assembling research teams, and procuring necessary equipment.

Let J_t denote the unit value of R&D ideas (i.e., ideas in ongoing R&D projects), and let $\Lambda_{t,t+1}$ represent the household's stochastic discount factor. Additionally, let $1-\phi_A$ denote the obsolescence rate of technology. The unit value of an idea within a pending R&D project is then expressed as:

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

For constructed R&D projects, the remainder of Phase II describes the process by which R&D ideas are converted into new varieties of intermediate inputs. Specifically, in each period, there is a probability λ that an R&D idea is adopted as a new variety.

To ensure an adoption rate of λ, the innovator must hire a fixed amount $\bar{n}$ of skilled labor. R&D ideas that are not adopted are subject to a depreciation rate of $1-\phi_A$.

To summarize, the unit value of R&D ideas, denoted by J_t, is:

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

where V_t denotes the unit value of a new variety, which will be determined in Phase III, and $\bar{n}$ is the fixed amount of skilled labor. The first term in (4.5) represents the cost of

skilled labor. The second term is the value of R&D ideas that are successfully converted into new varieties of intermediate inputs. The third term is the expected value of R&D ideas that remain to be adopted.

Phase III: Monetizing the Intellectual Property Rights Once adopted, an R&D idea within an R&D project transforms into a one-period patent (Acemoglu et al. 2012), allowing the innovator to sell the patent to an intermediate input producer and claim the profit from production for a single period. The value of the patent is defined as:

$$V_t = \Phi_t, \tag{4.6}$$

where Φ_t is defined in Equation (4.17) below, in the production block of the model.

The degree of cyclicality of innovation depends on both the value of a patent and the adoption cost $W_t^{RD}\bar{n}$ in (4.5), which together determine the unit value of R&D ideas that ultimately influences innovators' extensive-margin decisions on whether to set up a new project in Phase I.[8] The new varieties created in the innovation block contribute to the aggregate number of varieties A_t, whose evolution is described next.

The Evolution of A_t Recall that we differentiate raw ideas, denoted by $s_{i,t}$, from R&D ideas. Given the time to build introduced in Equation 4.4, we further distinguish between R&D ideas in a built project, denoted by $x_{i,t}$, and R&D ideas in a pending project, denoted by $\tilde{x}_{i,t}$.

Proposition 1 implies a relationship between $s_{i,t}$ and $\tilde{x}_t$. Specifically, the amount of new R&D ideas in a pending project, denoted by $\tilde{x}_{i,t}^{new}$, is:

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

Let $\tilde{X}_t \equiv \int_i \tilde{x}_{i,t}\, di$ denote the aggregate amount of R&D ideas in pending R&D projects at the end of period t. $\tilde{X}_t$ evolves as follows:

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

where $(1-\tilde{\lambda})\tilde{X}_t$ represents the R&D ideas in the pending projects that were not success-

[8] As discussed in Online Appendix C.1, the specific choice we make in (4.6) simplifies the model by avoiding the need to keep track of the idea distribution. Similar to Barlevy (2007), the parameter $\bar{n}$ can help to match any degree of cyclicality for a given procyclicality of V_t implied from the modeling choice of patent values. Therefore, the mechanism we emphasize does not rely on any specific assumption about patent values.

fully built in period t, and $\int_i \tilde{x}^{new}_{i,t+1}\, di$ captures the aggregate amount of new R&D ideas in pending projects.

Similarly, the law of motion for the aggregate R&D ideas in built projects, denoted by $X_{i,t}$, is:

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

Finally, the aggregate number of varieties, A_t, evolves according to:

$$A_{t+1} = \phi_A A_t + \lambda X_t A_{t+1}. \tag{4.10}$$

A_t is subject to a depreciation rate of $1-\phi_A$.

An important lesson that can be drawn from this section is that business cycle fluctuations affect the distribution of $s_{i,t}$ and $\bar{s}_{i,t}$, and therefore, impact $\tilde{x}^{new}_{i,t}$ through Equation (4.7). This, in turn, ultimately affects the aggregate productivity A_t through Equations (4.8) to (4.10).

4.2 The Production Block

The production block refers to the production of intermediate goods $Y_{t,m}$ and the *production* of different varieties of intermediate inputs $M_t(j)$. The latter differs from the *creation* of new varieties of intermediate inputs, which was discussed in Section 4.1.

Intermediate Goods Producer: There is a representative intermediate goods producer who produces and sells goods in a competitive market, using unskilled labor L_t and an aggregate intermediate input M_t. M_t is a composite of differentiated intermediate inputs $M_t(i)$:

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

where $\vartheta > 1$ is the parameter affecting the elasticity of substitution between intermediate inputs, and A_t is the total number of varieties of intermediate inputs. Following Romer (1990), the creation of new intermediate inputs, i.e., the expansion of A_t, is the source of productivity growth. The intermediate inputs $M_t(i)$ are created by innovators, as discussed in Section 4.1.

The intermediate goods producer faces two optimization problems. The first involves

minimizing costs when purchasing intermediate inputs, which leads to the downward-sloping demand function for $M_{t,j}$:

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

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 aggregate intermediate input. This demand equation is relevant for an intermediate input producer's profit optimization problem that we will discuss later.

The intermediate goods producer's second optimization problem involves the choice of inputs used in production, subject to the following production technology:

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

where $Y_{t,m}$ denotes intermediate goods, and γ is the parameter governing the share of the intermediate *inputs* in the production of intermediate *goods*. The second optimization problem can be described as the following profit maximization problem:

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

subject to production function (4.13). $P_{t,m}$ is the price of the intermediate good, and W_t is the wage of unskilled labor. The associated first-order conditions are:

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

Intermediate Input Producer Next, we turn to the production of the intermediate input $M_{t,j}$. The intermediate input producer, who claims the variety j of the intermediate input, produces $M_{t,j}$ with a unit cost of intermediate goods. In other words, the production function for the intermediate input producer j is given by:

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

The intermediate input producer operates in a monopolistically competitive market, facing a demand function characterized by (4.12). The profit-maximizing intermediate input

producer charges an optimal price that equals its desired markup (ϑ) times the marginal cost ($P_{t,m}$), $P_{t,j}^M = \vartheta P_{t,m}$. The resulting profit of the intermediate input producer j is denoted by $\Phi_{t,j}$:

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

Since the production process is symmetric across intermediate input producers, $\Phi_t = \Phi_{t,j}$ for all j. Hereafter, we drop the j index for intermediate input producers and work with Φ_t, denoting the profit of any intermediate input.

To recap, the production block connects with the innovation block in two dimensions. First, the profit of a new variety of intermediate input, Φ_t, as given by (4.17), serves as an incentive for innovation. Second, innovation is the fundamental driver of economic growth, as a new variety expands A_t according to (4.10), which in turn leads to growth according to (4.11) and (4.13).

4.3 Business Cycle Block

The business cycle block consists of a representative household, retailers, and a central bank.

The Household A representative household consumes and saves with risk-less bonds in zero supply. The representative household supplies two types of labor: unskilled labor L_t and skilled labor N_t. The unskilled labor is supplied to intermediate goods producers for producing intermediate goods, while the skilled labor is supplied to innovators for the adoption of R&D ideas into new varieties of intermediate inputs.

Denote C_t as consumption, Q_t as the holdings of risk-free bonds, W_t as the nominal wages of unskilled labor, W_t^{RD} as the nominal wages of skilled labor, and R_t as the return on risk-free bonds. Ξ_t collects the dividends paid by innovators, intermediate input producers, and retailers. The representative household's problem is given by

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

subject to

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

where ϱ_t is a demand shock given by the following exogenous process: $\log \varrho_{t+1} = \rho_\varrho \log \varrho_t + \sigma_\varrho e_{t+1}^\varrho$, with $e_{t+1}^\varrho \sim N(0,1)$. Solving the household's problem leads to the consumer Euler equation:

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

where $\Pi_{t+1} \equiv \frac{P_{t+1}}{P_t}$ is the inflation rate and $\Lambda_{t,t+1} = \beta \frac{C_t^\sigma}{C_{t+1}^\sigma} \frac{\varrho_{t+1}}{\varrho_t}$ is the stochastic discount factor. The first-order conditions for the two types of labor supply are as follows:

$$C_t^\sigma L_t^\varphi = \frac{W_t}{P_t}, \qquad C_t^\sigma N_t^\eta = \frac{W_t^{RD}}{P_t}. \tag{4.21}$$

Retailers and Final Goods We introduce nominal price rigidity following the "Rotemberg" type price adjustment cost framework.

Two types of agents play distinct roles in the process of transforming intermediate goods into final goods. Initially, retailers procure and convert intermediate goods into differentiated consumption goods, $y_{t,h}$; these are then sold to the final goods producer, who compiles the differentiated consumption goods into final goods denoted by Y_t.

The final good producer operates in a perfectly competitive market and transforms $y_{t,h}$ into the composite of the final good using the following function:

$$Y_t = \left(\int y_{t,h}^\mu \right)^{\frac{1}{\mu}} \tag{4.22}$$

where μ is a parameter affecting the substitution elasticity of the differentiated consumption goods. The corresponding price index of the composite final good is defined as $P_t = \left(\int p_{t,h}^{\frac{1}{1-\mu}} \right)^{1-\mu}$. The final good producer's cost minimization leads to the following demand function:

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

There is a fixed mass of retailers $h \in [0,1]$ who produce differentiated consumption

goods using the intermediate goods and sell in a monopolistically competitive market. The production technology is one-to-one, where one unit of intermediate goods is used to produce a unit of consumption goods:

$$y_{t,h} = Y_{t,m}. \tag{4.24}$$

Retailers set prices subject to quadratic price adjustment costs, given by $\frac{\kappa}{2}\left(\frac{p_{t,h}}{p_{t-1,h}} - 1\right)^2 Y_t$, where $p_{t,h}$ represents the price of consumption goods of the variety h, and Y_t is the composite of the final good.

The retailer h's optimization problem is to set price to maximize its profit:

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

subject to the demand constraint (4.23). $\Lambda_{t,t+j} \equiv \prod_{k=0}^{j-1} \Lambda_{t+k,t+k+1}$. Denoting $\pi_t \equiv \frac{P_t}{P_{t-1}} - 1$, and log-linearizing the optimality condition of the retailer's problem lead to the familiar New Keynesian Phillips Curve (NKPC):

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

Monetary Policy The central bank sets the nominal interest rate in response to productivity growth, $\frac{A_{t+1}}{A_t}$, and inflation, Π_t, according to the following Taylor rule:

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

A Taylor rule that responds to productivity growth is desirable, as the latter is proportional to the efficient real interest rate (Wicksell, 1898). More importantly, this assumption enhances tractability of the model, which we discuss next.

Discussions of Model Assumptions Online Appendix C.1 provides detailed discussions of model assumptions, including, but not limited to, the assumptions we made to make the model tractable. One aspect worth emphasizing is that, in the model, distribution matters for the evolution of aggregate productivity A_t, and therefore, distribution can affect households' consumption decisions. In this case, one would need to keep track of the distribution of $s_{i,t}$ and $x_{i,t}$. Since distributions are infinite dimensional variables,

employing numerical methods to handle them would significantly increase the model's complexity and reduce its interpretability.[9] The New Keynesian block, together with the assumed Taylor rule, ensures that the real interest rate does not depend on the growth rate of the economy, and therefore, R&D decisions and distributions do not affect households' consumption decisions. This feature greatly simplifies the model and maintains the focus of our analysis on how the distribution alters R&D decisions.

Online Appendix C.1 presents additional discussions on the model assumptions regarding the depreciation of ideas, passive adaptation, and the time-to-build.

4.4 Market Clearing

The market clearing condition for the final and intermediate goods markets are given by:

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

The labor market clearing requires that the demand for skilled/unskilled labor equals the supply of skilled/unskilled labor implied by the household problem. Finally, the market for risk-free bonds must clear, which implies that in equilibrium, risk-free bonds are in zero net supply $Q_t = 0$. Online Appendix C.2 defines the model's equilibrium, summarizes the equations characterizing it, and describes the solution method.

5 Model Calibration and Estimation

We calibrate the model to the U.S. economy at an annual frequency. Section 5.1 presents the calibration strategy, while Section 5.2 discusses the calibrated parameters and the model fit. In Section 5.3, the estimation of shocks is presented, given the calibrated model.

5.1 Calibration

We set standard parameters to values corresponding to those found in the literature. Parameters crucial for the quantitative results of the paper are internally calibrated to match selected moments in the data. These data moments include statistics related to firms' spike adjustments in R&D expenditures, which we now proceed to describe.

[9] Moreover, this complication arising from the distribution interacts with the kinked nature of innovators' extensive margin decisions. For instance, the approximation by Krusell and Smith (1998) generates substantial errors when applied to the current context, which features kinked policy functions.

Fact: R&D Spikes and Their Procyclicality In the model, an innovation decision along the extensive margin is associated with a spike adjustment in R&D expenditure. In the data, we define R&D spikes as increases in R&D expenditures that exceed a given threshold. To ensure the robustness of our findings to the choice of threshold, we consider a range of values from 40% to 100%. Specifically, we compute statistics by pooling R&D spikes identified under different thresholds. To avoid capturing trivial adjustments or structural changes, we exclude changes below $10,000 and those due to acquisitions exceeding 5% of total assets from the definition of a spike adjustment. Based on this definition, we identify, for each firm at each point in time, whether it has experienced a spike adjustment in R&D. The R&D spike *rate* is defined as the fraction of innovating firms (at time t) exhibiting spikes in their R&D expenditures.

Figure A.3 plots the evolution of the fraction of R&D spikes against real GDP growth rate across time. The correlation between these two series is 0.52 (p-value $= 0.00$), suggesting that firms usually choose to make significant changes in R&D expenses during economic booms.[10]

Fixed Parameters The discount factor β is set to match an annual interest rate of 4%. Parameter φ is set to 0.5, corresponding to the inverse Frisch elasticity for unskilled labor of 2. The intertemporal substitution parameter σ is set to 1, and the Frisch elasticity parameter for skilled labor, η, is set to 0 to ensure the tractability of the model discussed above.

For the New Keynesian block, we set $\kappa = 90$ and $\mu = 1.1$ to generate a Phillips Curve slope equal to 0.1, as suggested by Ottonello and Winberry (2020). The Taylor rule coefficient on inflation, φ_π, is set to 1.69 based on our own estimation of the Taylor rule. The persistence of demand shocks, ρ_ϱ, is set to 0.4, corresponding to a quarterly persistence parameter of 0.8.

For the parameters related to innovation and production, we follow Comin and Gertler (2006) to set the share of technology goods (intermediate inputs), γ, to 0.5. ϑ is set to 2, such that $\vartheta\gamma = 1$, following Kung and Schmid (2015). We set the parameter that governs the depreciation of intermediate inputs, ϕ_A, to 0.97, implying an annual depreciation rate of 3% as per Comin and Gertler (2006). The adoption rate, λ, is set to 0.2 to reproduce an average time of adoption of five years (Anzoategui et al. 2019). The elasticity of successful innovation with respect to R&D expenditures, ξ, is set to 0.5 following Acemoglu et al.

[10] In Online Appendix B.2, we provide a detailed discussion of additional robustness checks for this fact, including various alternative definitions of R&D spikes.

(2018).

Calibration Strategy Next, we discuss the strategy for the calibration of the key parameters: ζ, σ_ϱ, $\bar{n}$, c, ϕ, and $\tilde{\lambda}$. These parameters are calibrated jointly to match selected moments, summarized in Table 1. Now, we explain the rationale behind the selection of data moments.

As discussed in Section 4.1, the inflow of ideas is given by $\eta_{i,t+1} = \zeta N_{i,t}^{\xi} \cdot \varepsilon_{i,t+1}^{\eta}$, which depends on innovators' choice of $N_{i,t}$. Since idea inflows are inherently additive, we assume that $\varepsilon_{i,t+1}^{\eta}$ is drawn from a normal distribution $\mathcal{N}(1,1)$, truncated below at zero, ensuring that $\varepsilon_{i,t+1}^{\eta} \geq 0$. As the inflow of ideas is the source of economic growth, we incorporate the average growth rate of real GDP to inform the calibration of ζ.

The lumpiness of innovation is governed by the parameter c, where the fixed setup cost of a new R&D project is cA_t. We include the average *time duration between spikes* in the targeted moments to inform the calibration of c.

Next, we turn to the parameter that governs the obsolescence of unused raw ideas, $1-\phi$. It determines the difficulty of further accumulating raw ideas as the idea stock approaches the adjustment threshold. The larger the obsolescence rate, the less likely that an R&D project contains an amount of raw ideas much larger than the adjustment threshold $\bar{s}_t$. In other words, a larger $1-\phi$ will lead to a smaller spike in R&D expenses. Therefore, we include the ratio between the largest and the average R&D expenditure (max-to-mean ratio) to inform the calibration of $1-\phi$. The probability of finishing a pending R&D project, $\tilde{\lambda}$, is designed to reflect the empirical observation that the increase in spike rate due to a boom typically persists for several years. Therefore, we calibrate $\tilde{\lambda}$ by matching the first-order autocorrelation of the spike rate.

Finally, the parameters σ_ϱ and $\bar{n}$ determine the variance and persistence of the aggregate output. Specifically, σ_ϱ is the source of business cycle fluctuations in the model, while $\bar{n}$ influences the variance of TFP growth by amplifying the fluctuations in the profitability of new varieties. Consequently, the observed volatility in output growth is jointly informative for both σ_ϱ and $\bar{n}$. To further distinguish between these two parameters, note that they induce fluctuations with different persistence. The TFP growth dynamics exhibit higher persistence than the fluctuations of business-cycle components. The cyclical frequency of output growth is, therefore, a result of both TFP and business-cycle component persistence, with each component's influence weighted by its respective volatility. Thus, we add the AR(1) correlation of the growth rate to further inform the calibration of σ_ϱ and $\bar{n}$.

Calibrated Parameters The internally calibrated parameter values are detailed in Table 1. To discuss a few calibrated parameters, the obsolescence rate for unused raw ideas is 15%, which is higher than the rate for technologies already in use. The parameter $\tilde{\lambda} = 0.40$ indicates that, on average, an R&D project takes about 2.5 years from planning to completion. Lastly, $c = 0.001$ implies that the fixed adjustment cost of R&D is approximately 8% of the value of ideas embedded in R&D projects. Overall, our model fits well the targeted moments, as evidenced in the right panel of Table 1.

5.2 Calibration Results

Table 1: Calibration: Fitted Parameters

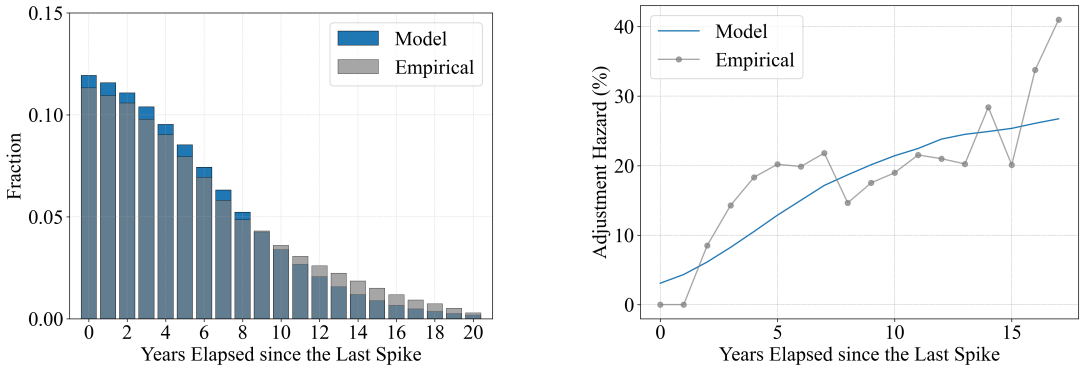
Parameter	Description	Value	Moments	Data	Model
			$E[\Delta \log(\text{GDP})]$	0.028	0.028
ζ	Innovation productivity	0.46	Time Duration between spikes	6.89	7.08
c	Fixed R&D cost	1e-3	Max-mean ratio R&D expense	2.40	2.41
ϕ	Idea obsolescence	0.15	AR(1) spike rate	0.57	0.51
$\tilde{\lambda}$	Pending R&D projects	0.40	AR(1) $\Delta \log(\text{GDP})$	0.36	0.37
$\bar{n}$	Skilled labor for adoption	0.79	$\sigma(\Delta \log(\text{GDP}))$	0.017	0.018
σ_ϱ	SD of preference shocks	0.02	Corr. Spike and $\Delta \log(\text{GDP})$	0.52	0.52
			Corr. R&D and $\log(\text{GDP})$	0.62	0.68

Notes: The sample period ranges from 1983 to 2019. The max-mean ratio of R&D expenses is computed as the average firm-level ratio of maximum R&D expenses to mean R&D expenses. Time duration between spikes is defined as the average number of years between two spike events for each firm.

Untargeted Moments: Procyclicality R&D Spike Rate In our internal calibration, we focused on matching micro-level moments that reflect the lumpy nature of innovation, such as the average duration between spikes and the mean-max ratio of R&D expenses. For untargeted moments, we concentrated on examining whether the model reproduces the aggregate characteristics of lumpy innovation. The last two rows of Table 1 report the empirical moments and the model's fit for the cyclicality of R&D and R&D spikes– neither of which is targeted in the calibration. The model closely matches the cyclicality of aggregate R&D. More importantly, it also captures the procyclicality of the R&D spike rate. We view this as the first validation of the model.

Untargeted Moments: Distribution Previously, we have emphasized the role of the distribution of raw ideas $s_{i,t}$. The credibility of the quantitative results presented in Section

Figure 5: Untargeted Moments: Distribution and Adjustment Hazard



(a) Years Elapsed since the Last Spike (b) Adjustment Hazard

Note: 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.

6, therefore, depends on the distribution of $s_{i,t}$ in the calibrated model. Is the distribution of $s_{i,t}$ realistic? Answering this question is challenging because ideas are not observable. In other words, there is no an empirical counterpart for the distribution of $s_{i,t}$.

The following proposition overcomes this challenge by showing that there exists a one-to-one mapping from the distribution of unobserved ideas to the distribution of an observable: the time elapsed since the last spike. Intuitively, Proposition 3 holds because, in our model, an innovator who has waited a long time since the last spike adjustment in R&D expenses is more likely to have a big $s_{i,t}$.[11]

Proposition 3. *Let $\mathcal{G}_t$ denote the cumulative distribution function (CDF) of $s_{i,t}$, and let $\mathcal{T}_t$ denote the CDF of the time elapsed since the last spike, represented by $TE_{i,t}$. For $n \in \mathcal{N}$, there exists a conditional distribution function of $s_{i,t}$, $P(\cdot|TE_{i,t} = n)$, which describes the conditional distribution of $s_{i,t}$ given $TE_{i,t} = n$. Based on this conditional distribution function, the following relationship holds:*

$$\mathcal{G}_t(s) = \sum_{n \in \mathcal{N}} P(s|TE_{i,t} = n)\mathcal{T}_t(n). \tag{5.1}$$

[11]The proof is straightforward and follows from the law of total probability and the definition of conditional probability.

Figure 5a presents both the empirical distribution of $TE_{i,t}$ and its model counterpart. Our model closely matches the overall shape of the $TE_{i,t}$ distribution. The fit of these untargeted micro-level moments allows us to make quantitative statements using the model and guide policymaking with confidence.

To further assess the empirical relevance of Proposition 3, we compute the hazard function using firm-level data and compare it to its theoretical counterpart. Figure 5b plots the empirical hazard function (grey line with circles) alongside the model-implied hazard function (blue solid line). The two curves closely overlap in both slope and level, providing strong support for the model's prediction.

5.3 Fitting the Model to the Data over Time

Given the calibrated parameters, we then proceed to the estimation of shocks. The estimated shocks, particularly those that triggered the Great Recession, will be used in the quantitative analysis presented in Section 6. We fitted our model to the U.S. GDP growth rate from 1984 to 2018 to estimate the time series of demand shocks, ε_t^Q.[12]

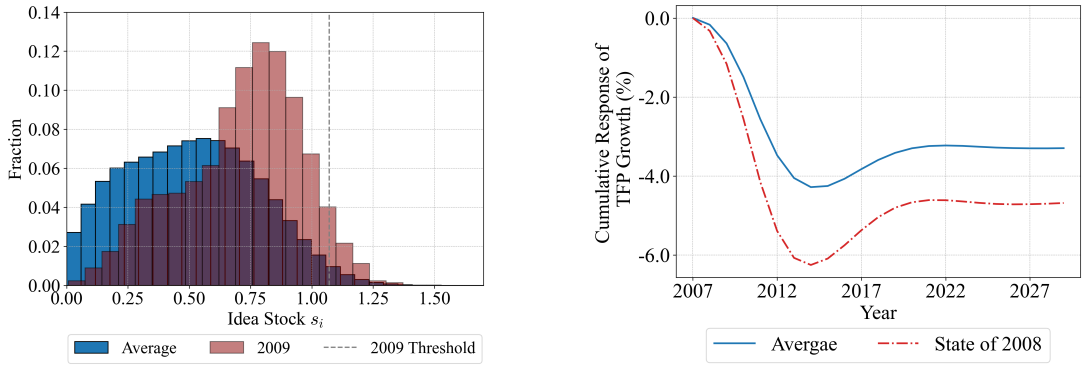
Figure A.4a shows the fitting results, where the model-derived GDP growth rate closely matches the actual rate. Additionally, Figure A.4a includes 80% and 99% probability bands for one-period ahead predictions, indicating that the GDP growth rate for period $t+1$ is likely to fall within this range based on the state at period t. As an external validation, Figure A.4b compares the model-implied growth of A_t, inferred from fitting the model, to U.S. GDP growth rate data, with the TFP growth rate computed based on Fernald (2014)'s capital utilization-adjusted TFP measure. Our model-implied growth of A_t closely aligns with the general trend of TFP growth in the data[13].

Our model suggests a gradual decline in TFP growth rate since the early 21st century, with the 2008 Great Recession merely continuing this trend. Existing literature typically attributes this decline to factors like slowed technological change in information and communications technology (Cette et al. 2016), slowed technology diffusion (Anzoategui et al. 2019) and the increasing difficulty of finding new ideas (Bloom et al. 2020). However, our model offers an alternative explanation, suggesting that the late 1990s dot-com bubble

[12]Specifically, we began the fitting from 1979 because the first five periods before 1984 are considered a burn-in phase to transition the model economy to the 1984 state. To avoid over-fitting, we minimize the mean square error plus a penalty term that is inversely proportional to the likelihood of the fitted shocks.

[13]The model-implied TFP growth is smoother. This is primarily because, in our model, TFP is solely influenced by technological progress, while actual aggregate TFP swings are more complex, often including factors like misallocation (e.g. Hsieh and Klenow 2009) and changes in labor force composition (Moscarini and Postel-Vinay 2018). Moreover, the estimated TFP in the data is inherently noisy.

Figure 6: Great Recession: Distribution Matters



(a) Distribution before the Great Recession (b) CIR of TFP to the Great Recession Shocks

Notes: Panel (a) compares the estimated distribution of $s_{i,t}$ in 2009 to the average cross-sectional distribution of $s_{i,t}$ implied by the model. Panel (b) displays the cumulative TFP response to the 2008 and 2009 shocks, computed by contrasting the actual trajectory with a counterfactual in which the 2008–2009 shocks are set to zero. The red curve shows the cumulative TFP response starting from the 2008 state, while the blue curve depicts the response computed under the average distribution of $s_{i,t}$.

significantly boosted R&D, which exhausted innovators' idea stock. The subsequent period of moderation in the early 21st century led the R&D sector of the economy into a phase of idea accumulation. During this period, the economy was primarily processing the stockpile of R&D built up during the dot-com bubble, which in turn contributed to the decline in TFP growth. We will explore this in greater detail in the next section.

6 Model Results

The main insight of this paper is to provide a theory of lumpy innovation and illustrate its implications for aggregate dynamics, with an emphasis on the role of the raw idea distribution. This section presents these results based on the estimated model. Section 6.1 demonstrates that the *shape* of the idea distribution contributed to the observed severity of the Great Recession. Section 6.2 shows that the preceding Great Moderation led to the observed shape of the idea distribution prior to the Great Recession. Finally, Section 6.3 demonstrates that moderation can hinder economic growth.

6.1 Great Recession: Distribution Matters

The distribution of idea stock that we emphasize contributed to the severity of the 2008 Great Financial Crisis. Hereafter, we refer to it as the GFC or the Great Recession interchangeably. To illustrate this, we retrieve the distribution of $s_{i,t}$ prior to the 2009 shock.[14]

The model features two state variables: the aggregate demand shock ϱ_t and the distribution of idea stock across firms. Figure 6a plots the distribution of the idea stock in 2009, according to the estimated model, alongside the average distribution. We observe that the 2009 pre-shock distribution, compared to the average, has a substantial mass concentrated around the adjustment threshold, implying a considerable risk of loss in R&D activities.

Figure 6b illustrates the decline in A_t resulting from the 2008 and 2009 recessionary shocks. The effects of the Great Recession shocks on A_t unfolded gradually.[15] In the long run (from 2030 onward), the Cumulative Impulse Response (CIR) of A_t indicates that the U.S. economy, having experienced the Great Recession, now exhibits a GDP nearly 4.7% lower than that of a counterfactual U.S. economy without the Great Recession shocks (see the red dashed line).[16] Such results are not surprising in a framework where business cycle shocks affect innovation.

Surprisingly, if the Great Recession shocks had occurred in an average state (with an average distribution), the loss in aggregate output caused by these shocks would have been more moderate, at around 3.3% in the long run (see the blue solid line). This implies that about 40 percent of the total economic loss can be attributed to the distribution of idea stock prior to the Great Recession.

6.2 The Great Twins: Great Moderation and Great Recession

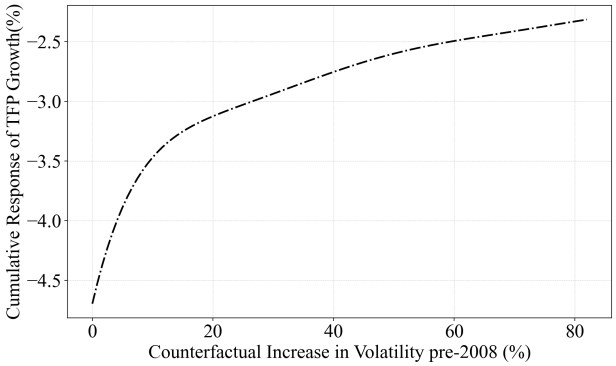
Next, we pose the question: What contributed to the 2009 pre-shock distribution of the idea stock, which displays a significant mass around the extensive margin? We argue that this pre-shock distribution–implying elevated left-tail risk–was a consequence of the moderated economic fluctuations during the 2002–2007 period.

[14]Since our fit is based on annual U.S. GDP growth data, the largest negative shock actually occurred in 2009.

[15]The average distribution of ideas (blue), plotted in Figure 6a, is simulated to incorporate the same $\hat{\varrho}_{t-1}$ in 2008 as in the realized economy (red). This ensures that the counterfactual analysis in Figure 6b is conditional on $\hat{\varrho}_t$ of identical magnitude: $\hat{\varrho}_t = \rho\hat{\varrho}_{t-1} + \varepsilon_t^{\varrho}$.

[16]Note that this does not imply GDP is 4.7% lower than its pre-shock level, since the counterfactual economy would have experienced positive growth.

Figure 7: The Great Twins



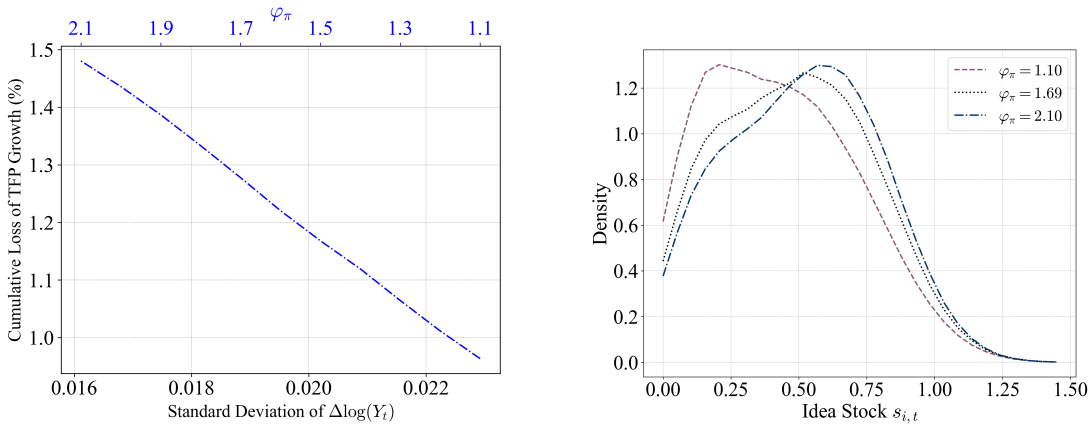
Notes: This figure plots the cumulative effects (40 years) of the 2008 and 2009 recessionary shocks on TFP growth across alternative counterfactual economies with different realized pre-shock sample volatilities. To this end, we compute the loss in TFP growth due to the 2008 and 2009 shocks by introducing additional exogenous shocks to the fitted values of ε_t^{ϱ} from 2003 to 2007. Specifically, the counterfactual shocks are defined as $\tilde{\varepsilon}_t^{\varrho} = \varepsilon_t^{\varrho} + \tilde{\sigma}\tilde{\varepsilon}_t$, where $\tilde{\varepsilon}_t$ is drawn from a standard normal distribution. For each value of $\tilde{\sigma}$, we compute both the average loss in TFP growth (y-axis) resulting from the 2008 and 2009 shocks and the corresponding average increase in the volatility of ϱ_t over the 2003–2007 period (x-axis).

To investigate this, we simulate a counterfactual U.S. economy over 2002–2007 by adding additional disturbances to the estimated sequence of aggregate demand shocks. Conditional on these added shocks, we compute the counterfactual increase in *sample* volatility over this period, which is plotted on the x-axis of Figure 7. For each counterfactual state of the economy, we then compute the associated loss in cumulative TFP growth resulting from the *same* 2008 and 2009 recessionary shocks. These losses are plotted on the y-axis of Figure 7. We find that increasing volatility in 2002–2007 by 20% reduces the average loss in A_t caused by the Great Recession shocks by 1.5 percentage points. That is, the moderation period following the dot-com bubble led to an over-accumulation of ideas, making the economy more vulnerable to the Great Recession shocks.

It is worth clarifying that reduced *sample* volatility, which we refer to as the Great Moderation, does not imply reduced risk. In the previous simulations, the moderation observed prior to 2008 results from realizations of shocks of different sizes, while the underlying risks–i.e., the *unconditional* or *population* volatilities of shocks and the economy–remain constant. This way of generating moderation reflects the "Good Luck" view of the Great Moderation, as emphasized by Justiniano and Primiceri (2008).

Policy Implication: The Tail-Risk Variance Trade-off An alternative view is that the Great Moderation is a result of a 'Good Policy' (Galí and Gambetti 2009): the U.S. Federal

Figure 8: Policy Implications of the Great Twins



(a) The Tail-Risk Variance Trade-off

(b) Distribution of $s_{i,t}$

Notes: Panel (a) illustrate the cumulative impulse responses to a severe adverse shock, across different policy parameter settings. In particular, the severe adverse shock is defined as a one-sigma negative preference shock that cannot be offset by monetary policy (i.e., $\varphi_\pi = 0$ throughout the recessionary period), effectively mimicking the constraints imposed by the zero lower bound. The top x-axis denotes the policy parameter values, while the bottom x-axis corresponds to the standard deviation of GDP growth. The cumulative loss in TFP is measured as the cumulative impulse response over 40 periods following the shock. Panel (b) plots the stationary distribution of $s_{i,t}$ across different policy parameter settings.

Reserve adopted a more effective monetary policy to mitigate excessive economic fluctuations, specifically, a larger ϕ_π in the Taylor rule. The Great Twins result holds regardless of the source of the Great Moderation.

Our model suggests that such 'Good Policy' has an unexpected negative consequence: increased tail risk. The central bank, by reacting more aggressively to inflation, can reduce unconditional volatility; however, conditional on the occurrence of a recession, the recession will be deeper and last longer. Figure 8a illustrates this result by plotting the Tail-Risk Variance trade-off induced by different ϕ_π.

The figure demonstrates that lowering φ_π increases overall economic volatility, as shown by the correspondence between the upper and lower horizontal axes. The vertical axis represents left-tail risk–specifically, the cumulative loss in TFP resulting from a one standard deviation negative demand shock that cannot be offset by monetary policy (i.e., $\varphi_\pi = 0$ throughout the recessionary period), effectively mimicking the constraints imposed by the zero lower bound. The downward-sloping blue dashed line indicates that reducing overall economic volatility through monetary policy adjustments can increase economic tail risk. This result is driven by differences in the shape of the idea distribution

generated under alternative monetary policy regimes. Figure 8b shows that an economy with a higher φ_π—and thus greater moderation—features an idea distribution with a fatter right tail compared to an economy with a lower φ_π.

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 40 years. How does this characteristic affect our results? In Online Appendix C.5, we show that for a lower average inflow of ideas, the 2008 Great Recession caused a greater loss of real GDP. This suggests an increasing importance of the mechanism/insight emphasized in our paper moving forward.

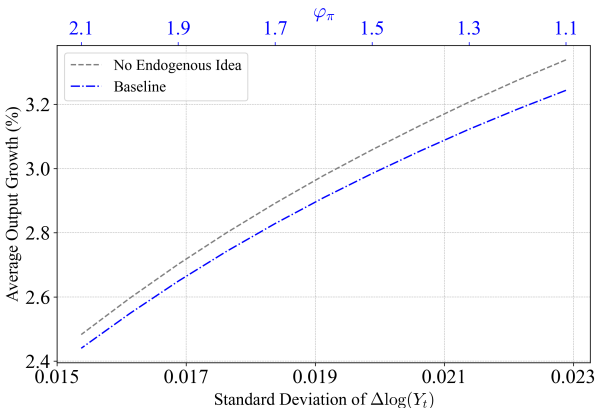
6.3 Great Moderation and Growth

The moderation-induced fat-tailed distribution also has implications for long-term growth. Specifically, an economy experiencing greater moderation–e.g., due to the conduct of monetary policy (larger φ_π)–tends to overaccumulate (or underutilize) raw ideas, as illustrated in Figure 8b. Under the assumption that unused raw ideas are subject to depreciation, this underutilization leads to losses in technological progress. As a result, moderation-induced underutilization of ideas is associated with lower average GDP growth. This result is shown by the blue line in Figure 9, where a higher value of φ_π (upper horizontal axis) is associated with lower volatility (lower horizontal axis) but also with lower average output growth (vertical axis).

In the baseline model, fluctuations affect growth both through their impact on the distribution–a channel we emphasize–and through their impact on idea creation. To isolate the former from the latter, the dashed black line in Figure 9 shows that a similar result holds in a version of the model with exogenous idea inflows.

These results offer a new perspective on the relationship between volatility and long-term growth, a topic that has been debated in both theoretical and empirical literature. On the theoretical side, the Schumpeterian view suggests that business cycles can promote long-term growth through creative destruction, whereby recessions accelerate the replacement of old technologies with new ones. In contrast, other models emphasize the negative effects of volatility on growth by increasing uncertainty and discouraging innovation (Jones et al., 2000; Aghion et al., 2010). These mechanisms are distinct from the moderation-induced fat-tailed distribution mechanism emphasized in our paper. Empirically, the direct relationship between growth and volatility is mixed (Kormendi and

Figure 9: Great Moderation and Growth



Notes: The blue dashed line plots the average output growth rate across different values of the monetary policy parameter φ_π. The grey dot-dash line plots the same relationship in a counterfactual setting that shuts down the endogenous growth channel by imposing a fixed idea inflow equal to the average idea creation across all states in the baseline model.

Meguire, 1985; Ramey and Ramey, 1995; Judson and Orphanides, 1999). More recent meta-analyses suggest that the relationship tends to be more negative in developing countries, and positive in advanced economies (Bakas et al., 2019).

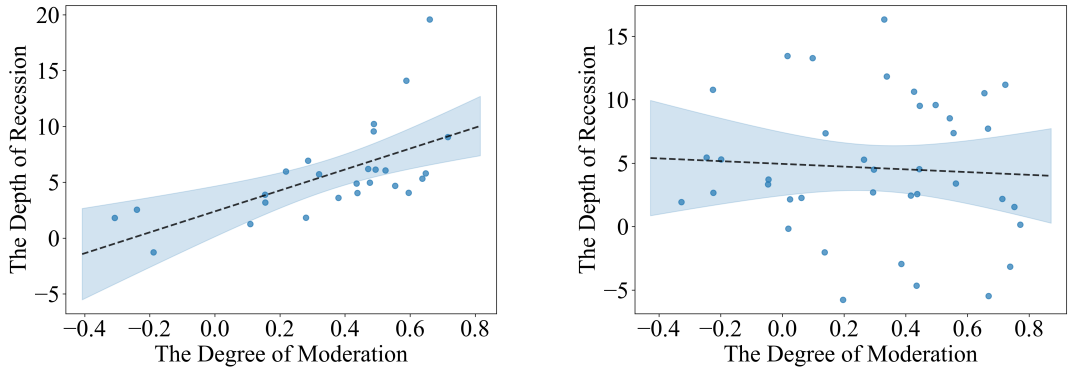
7 Empirical Validations of the Theory

Next, we return to the data to seek evidence that supports the model's predictions and its underlying mechanisms. Section 7.1 presents the Great Twins in the data: the positive correlation between the depth of (Great) recessions and the degree of moderation in the preceding periods. The key mechanism that produces the prolonged Great recession in our model is the distribution of ideas. Section 7.2 presents micro-level evidence that supports this mechanism. Section 7.3 demonstrates the relevance of extensive-margin R&D decisions for firm-level innovation outcomes.

7.1 The Great Twins in the Data

Turning to the macro-level evidence, we first focus on the Great Recession. Ball (2014) finds that the losses in output following the 2008 recession range from negligible sizes for countries like Australia and Switzerland to more than 30% in Greece, Hungary, and Ireland. Can the heterogeneity in losses be explained by the preceding degree of moder-

Figure 10: The Great Twins in the Data



(a) R&D-to-GDP $> 1\%$ (b) R&D-to-GDP $< 1\%$

Notes: The result is based on the cross-country data from the World Bank. To avoid the depth of recession being influenced by the pre-recession trend, we normalize the depth of recession by the pre-recession trend growth, defined as the average growth rate over the past fifteen years. The correlation for the countries with R&D-to-GDP ratio greater than 1% (Panel a) is 0.61 ($p = 0.00$) and for the countries with R&D-to-GDP ratio smaller than 1% (Panel b) is -0.06 ($p = 0.70$).

ation? We investigate the plausibility of the Great Twins hypothesis in the data.

To measure the degree of moderation, we calculate a country's standard deviation of GDP growth in the subsample before the recession and divide it by its full-sample standard deviation. The depth of the recession is defined by the gap between the actual GDP and the GDP predicted by the pre-recession growth trend.[17]

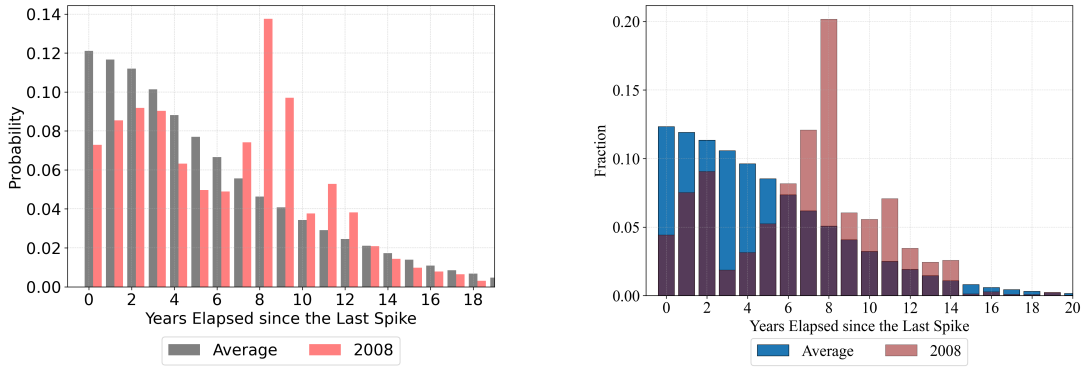
Figure 1 in the introduction presents the main finding: the depth of recessions is positively correlated with the degree of moderation in the pre-crisis periods across OECD countries, consistent with the Great Twins hypothesis. Figure 10a demonstrates this relationship for countries with an R&D-to-GDP ratio greater than one percent[18].

Our theory describes advanced economies that conduct innovation; it does not fit developing economies or economies that do not conduct R&D. Consistent with this prediction, we find that for countries with an R&D-to-GDP ratio below 1%, the correlation between the degree of moderation and recession depth disappears: -0.06 ($p = 0.70$). See Online Appendix B.3 for a more detailed robustness check of the sample selection rule.

[17] See Online Appendix B.3 for details of the formal definition of the depth of recession.

[18] We select countries with an R&D-to-GDP ratio greater than one percent because a low ratio may indicate that growth is not primarily driven by technological advancements, making the Great Twins hypothesis less applicable. The average R&D-to-GDP ratio for OECD countries is around 2%.

Figure 11: Distribution of Time Elapsed since the Last Spike: Pre-2008 Recession.



(a) Distribution: in the Data

(b) Distribution: in the Model

Notes: Panel (a) plots the distribution of time elapsed since the last spike in the data. To avoid outliers, the data of time elapsed since the last spike is truncated below 20 years. Panel (b) depicts the simulated distribution of time elapsed since the last spike and the distribution of time elapsed since the last spike retrieved from fitting model to the actual U.S. GDP growth.

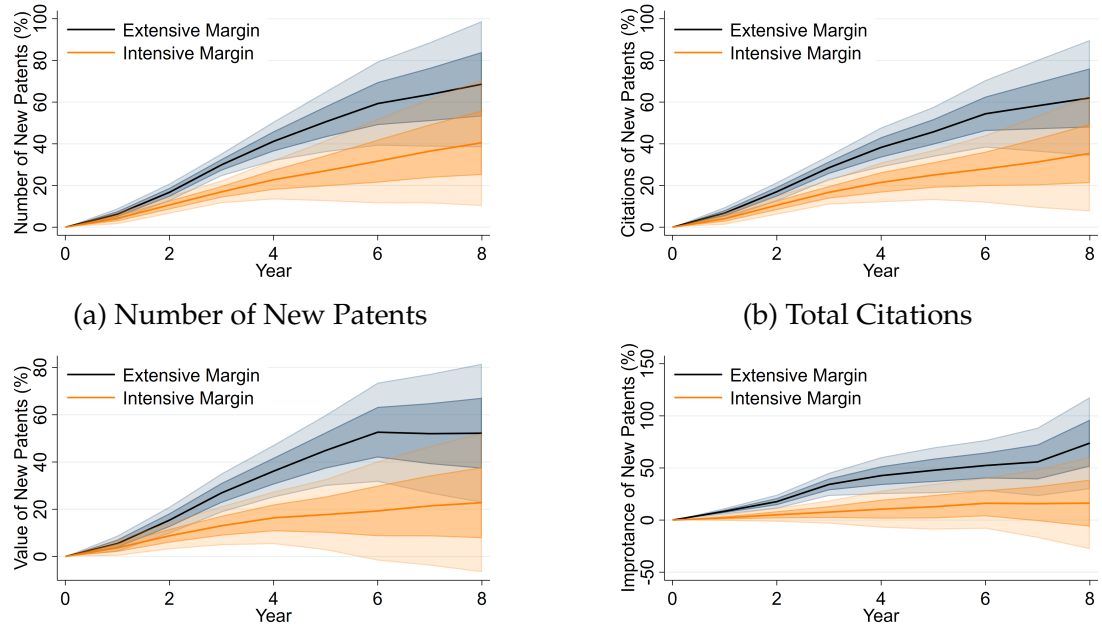
Robustness Checks In Online Appendix B.3, we demonstrate that the Great Twins relationship persists even when controlling for financial factors. Additionally, we find that the model-implied relationship between recessions and their preceding moderation is evident when considering TFP growth (instead of real GDP) and across other recessions (beyond the 2008 recession), using data from 18 advanced economies since World War II. We conclude that the Great Twins is a robust phenomenon across advanced economies and over time.

7.2 Micro-Evidence: Distribution in the Data

We now turn to the micro-level data to seek evidence supporting the mechanism. Specifically, the model suggests that the idea stock distribution, as indicated in Figure 6a, is crucial for the observed depth of the Great Recession. In Section 5.1, we discussed a challenge in verifying the mechanism: the idea stock distribution is a key state variable in the model, but its empirical counterpart is unmeasurable in reality. As proposed in Proposition 3, the solution to this challenge is to estimate the distribution of the time elapsed since the last spike, which has a one-to-one correspondence with the idea stock distribution.

Figure 11 compares the empirical and model-reconstructed histograms of the time elapsed since the last spike in 2008. The model-generated histogram suggests that the concentration of mass near the extensive margin, shown in Figure 11b, is due to a large

Figure 12: Lumpy Innovation and Patenting



(a) Number of New Patents

(b) Total Citations

(c) Total Value

(d) Total Importance

Note: The blue curve shows the cumulative impulse response of firms' patenting activity following an R&D spike (defined as an increase greater than 40%), while the orange curve reflects the response to a 100% increase in R&D expenditure. Shaded areas denote 95% and 68% confidence intervals, computed using Driscoll and Kraay standard errors.

number of innovators being 5-8 years past their last R&D spike, coinciding with the average 6.05-year gap between R&D spikes in the data. As shown in Figure 11a, this pattern is also mirrored in the empirical data!

7.3 Lumpy Innovation and Patenting

By emphasizing the importance of the distribution, our theory highlights innovators' decisions along the extensive margin. There are two testable implications along this dimension. First, innovators' extensive-margin decisions are procyclical. This evidence was presented in Section 5.1; see Figure A.3. Second, innovation along the extensive margin has real impacts on innovation outcomes–particularly, patents. This section confirms this prediction using firm-level data.

To test the second prediction we use firm-level patent data linked from the USPTO to Compustat, following the matching procedure in Arora et al. (2021). In addition to examining the quantity of innovation, as measured by the number of patents, we con-

sider three distinct measures of patent importance. The first is the number of citations received by each patent. The second is the patent value measure from Kogan et al. (2017), which captures the stock market reaction to patent grant announcements and reflects the economic value attributed to the innovation by investors. The third is the importance measure proposed by Kelly et al. (2021), which is constructed using natural language processing methods to compute the ratio of forward similarity to backward similarity of patent texts, thereby identifying patents that are both novel and influential.

We estimate the impact of R&D spikes on subsequent innovation outcomes using the following panel regression model:

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

where $y_{i,t+h}$ denotes outcomes of innovation—both quantity and quality measures. To mitigate the sparsity of patent output, we smooth each outcome using a three-year backward moving average, and compute the dependent variable $\Delta_h y_{i,t+h}$ as a forward percentage change relative to the average of the initial and final values across horizon h, following the standard practice in the patent innovation literature.

The variable $\text{Spike}_{i,t}$ is a dummy that equals one when firm i increases its R&D spending by more than 40 percent in year t. The coefficient β_h thus captures the effect of extensive-margin adjustments in R&D on future innovation outcomes. In contrast, the coefficient γ_h identifies the intensive-margin effect, based on the continuous change in R&D spending. The vector $X_{i,t}$ includes lagged firm-level controls such as sales, total assets, and R&D expenditure.

Figure 12 displays the cumulative response of innovation outcomes to changes along both the extensive and intensive margins of R&D. The blue curve corresponds to the effect of an extensive-margin adjustment, captured by an R&D spike, while the orange curve represents the impact of a 100% increase in R&D spending along the intensive margin.

The figure reveals that extensive-margin adjustments in R&D expenditures have economically significant effects on patent outcomes. Specifically, the effect of an R&D spike is comparable to that of a 200% increase in R&D expenditure along the intensive margin. This pattern holds across both quantity and quality measures, suggesting that extensive-margin decisions enhance not only the volume of innovation but also its quality.

8 Conclusion

We propose a theory of lumpy innovation and study its implications for aggregate dynamics. Specifically, we develop a business cycle model with heterogeneous innovators who face a fixed setup cost when initiating new R&D projects, giving rise to innovation decisions along the extensive margin. We highlight a novel mechanism–moderation-induced fat-tailed distributions–arising from the asymmetric effects of business cycle fluctuations on the distribution of ideas. This mechanism gives rise to the "Great Twins" hypothesis: the Great Moderation caused the observed severity of the subsequent Great Recession. Moreover, it leads to the underutilization of ideas, which can hinder long-run economic growth.

To support this new theory, we document the Great Twins in the data, uncover micro-level evidence consistent with the underlying mechanism, and show that extensive adjustments in R&D expenditures have significant effects on innovation outcomes, thereby validating the importance of the extensive margin emphasized in the model.

Overall, we view the main contribution of our paper as providing a new qualitative insight: moderation can causally contribute to a more severe recession. Quantitatively, according to our calibration, the Great Recession was 40% more severe due to the preceding Great Moderation. However, for those interested in understanding the full set of contributing factors to the Great Recession, a more thorough quantitative analysis that compares our mechanism with alternative explanations would be a fruitful avenue for future research. Additionally, future work could explore whether the mechanism we emphasize–arising naturally from lumpy innovation with one-sided adjustment–also manifests in other economic environments. Lastly, we adopt a Romer-type model primarily due to its tractability in environments with aggregate shocks. A similar modeling strategy—where firms accumulate ideas and face discrete decisions to initiate innovation projects—can also be applied to product-line-based models such as Klette and Kortum (2004). In such settings, idea accumulation would occur at the firm level, and R&D investment decisions would determine the timing of innovations that improve existing product lines or introduce new ones. We expect our findings to apply in these settings as well. Future research that incorporates idea accumulation and extensive-margin innovation decisions into product-line-based models such as Klette and Kortum (2004) could provide richer micro-level R&D dynamics and allow exploration of their aggregate consequences within a multi-product firm environment.

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 and Peter Howitt, "A Model of Growth Through Creative Destruction," *Econometrica*, 1992.

_ , 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.

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

Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez, "Endogenous technology adoption and R&D as sources of business cycle persistence," *American Economic Journal: Macroeconomics*, 2019, *11* (3), 67–110.

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 cyclicality of research and development," *American Economic Review*, 2007, *97* (4), 1131–1164.

Benigno, Gianluca and Luca Fornaro, "Stagnation traps," *The Review of Economic Studies*, 2018, *85* (3), 1425–1470.

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

Bilbiie, Florin O and Marc J Melitz, "Aggregate-demand amplification of supply disruptions: The entry-exit multiplier," Technical Report, National Bureau of Economic Research 2020.

_ , 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.

_ , _ , John C Haltiwanger, Michael Woodford, and Robert E Hall, "Plant-level adjustment and aggregate investment dynamics," *Brookings papers on economic activity*, 1995, pp. 1–54.

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.

Chiavari, Andrea and Sampreet Singh Goraya, "The Rise of Intangible Capital and the Macroeconomic Implications," *Working Paper*, 2024.

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.

Fabrizio, Kira R and Ulya Tsolmon, "An empirical examination of the procyclicality of R&D investment and innovation," *Review of Economics and Statistics*, 2014, *96* (4), 662–675.

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.

Gnewuch, Matthias and Donghai Zhang, "Monetary policy, firm heterogeneity, and the distribution of investment rates," *Journal of Monetary Economics*, 2025, *149*, 103721.

Gornemann, Nils, Pablo Guerrón-Quintana, and Felipe Saffie, "Real Exchange Rates and Endogenous Productivity," 2021.

Griliches, Ziv, "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 1990, *28* (4), 1661–1707.

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," *Econometrica*, 2008, *76* (2), 395–436.

Klette, Tor Jakob and Samuel Kortum, "Innovating firms and aggregate innovation," *Journal of political economy*, 2004, *112* (5), 986–1018.

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.

Lippi, Francesco and Aleksei Oskolkov, "A Structural Model of Asymmetric Lumpy Investment," 2023.

Matsuyama, Kiminori, "Growing through cycles," *Econometrica*, 1999, *67* (2), 335–347.

Miao, Jianjun and Pengfei Wang, "A Q-theory model with lumpy investment," *Economic theory*, 2014, *57*, 133–159.

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.

Ouyang, Min, "On the Cyclicality 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.

Wicksell, Knut, "Interest and Prices (English translation by RF Kahn, 1936)," 1898.

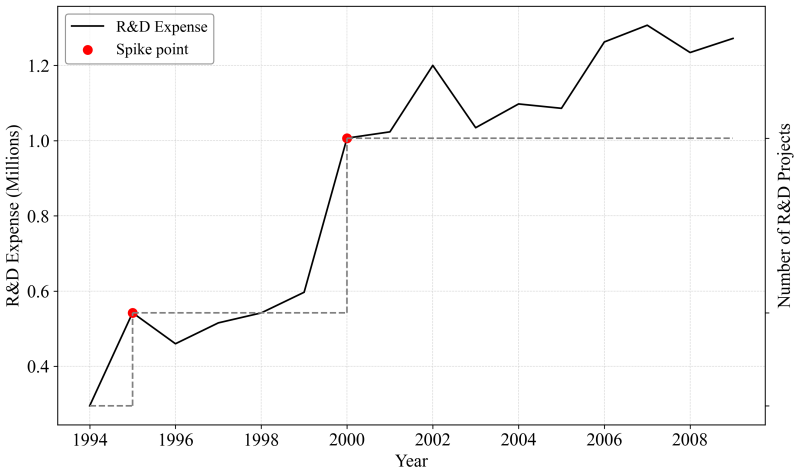
Winberry, Thomas, "Lumpy investment, business cycles, and stimulus policy," *American Economic Review*, 2021, *111* (1), 364–96.

Online Appendix to "A Theory of Lumpy Innovation and Aggregate Dynamics´´

Lixing Wang *Donghai Zhang*

A Additional Figures

Figure A.1: The Concept of Lumpy Innovation



Notes: This figure plots the time series of the R&D expense of a listed firm.

Figure A.2: Block Diagram of the Model System – the Innovation Block

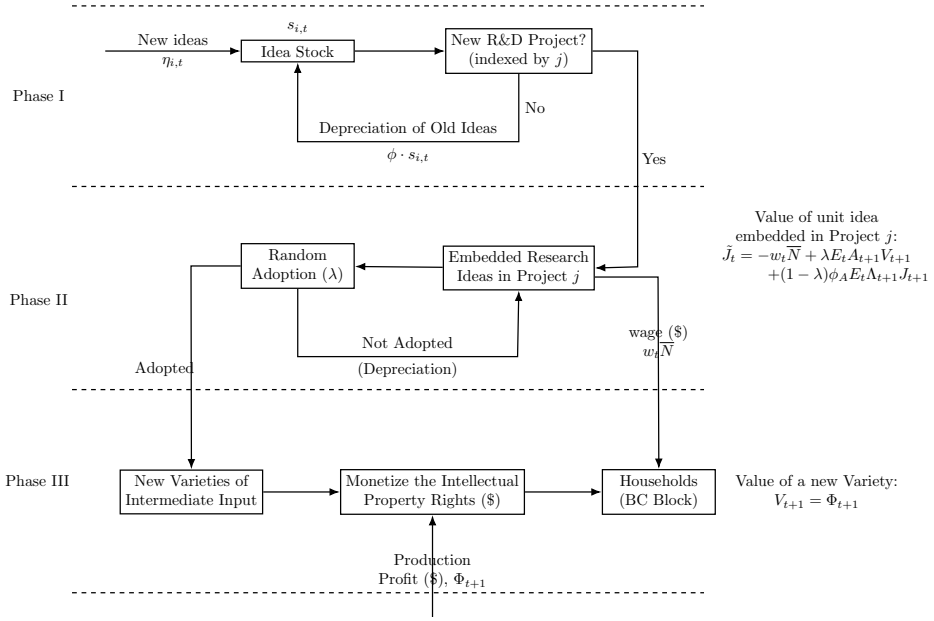
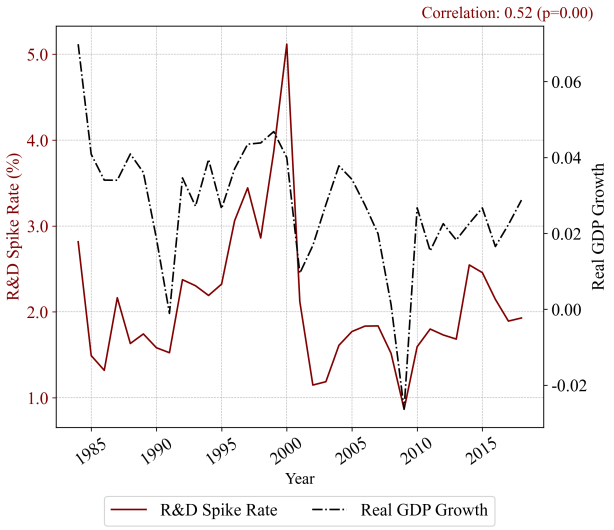
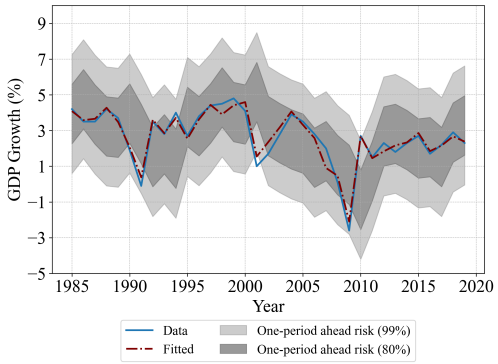


Figure A.3: The Pro-cyclicality of R&D spikes.

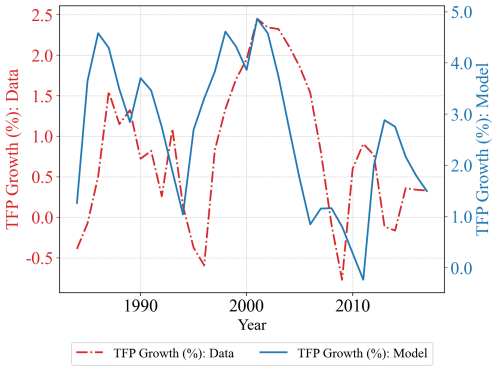


Notes: This figure plots the time series of the R&D spike rate, defined as the fraction of R&D spikes, along with real GDP from 1984 to 2018.

Figure A.4: Fitting Model to the Data over Time.



(a) Model Fit (in-sample): GDP



(b) Model Fit (out-of-sample): TFP

Note: The TFP data is from the Federal Reserve Bank of San Francisco, which reports updated capital utilization-adjusted TFP using Fernald (2014)'s method. Because the TFP series contains other factors beyond technology growth (e.g., capital misallocation), the red line plots a three-period moving average of TFP.

B Data Appendix

B.1 Sample Selection

We use the Compustat Fundamentals Annually database. Observations are uniquely identified by GVKEY and FYEAR. In line with the literature, we exclude observations which fall under the following criteria:

1. not incorporated in the United States
2. not reported in U.S. Dollar
3. in the following sectors:
    - Utilities (SIC 4900-4999)
    - Financial Industry (SIC 6000-6999)
    - Non-operating Establishments (SIC 9995)
    - Industrial Conglomerates (SIC 9997)
4. acquisitions exceed 5% of total assets
5. firms with less than five R&D expense observations
6. firms with no positive R&D expense observations
7. negative R&D expense

Table B.1 presents industry statistics for our selected sample and the full sample. The table highlights that our R&D firms are well-distributed across various industries. This comparison demonstrates that our sample selection does not overly focus on any single sector, ensuring the robustness and general applicability of our findings.

B.2 Lumpy Innovation: Discussion and Robustness Checks

Tax Incentive One of the concerns regarding the result of Fact 1 could be that the lumpy pattern might be a result of strategic accounting for tax reasons. We argue that this is not the case in the data. According to the rule of Credit For Increasing Research Activities in the U.S., the R&D Tax Credit provides a fixed percentage of the taxpayers qualified

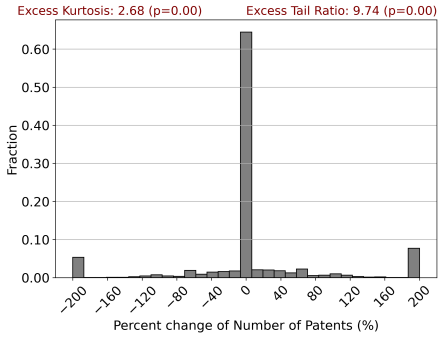
Table B.1: Industry Statistics

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

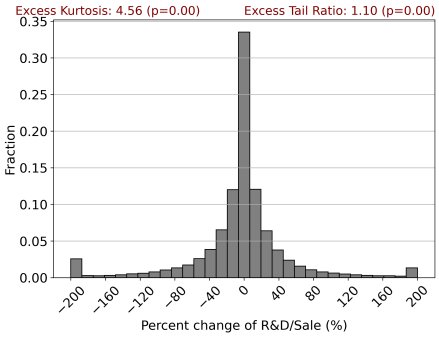
research expenses (QRE) that exceed a calculated base amount, which depends on the calculation methods. The design of the credit calculation incorporates features that encourage consistent/smoothed research spending over time due to the following reasons. First, the base amount for calculating the credit is tied to a company's historical QREs. Furthermore, regardless of the chosen calculation method, the base amount is always at least 50% of the taxpayer's current year qualified R&D expenses. As a result, the tax codes offer minimal incentives for strategically manipulating R&D expenditures such that they appear in lumpy patterns.

Lumpy Patents Lumpy innovation might imply that the outputs of R&D, such as patents, should also exhibit a lumpy pattern. Figure B.1a presents the distribution of changes in new patent counts for Compustat firms, based on USPTO patent data matched to Compustat following the approach in Arora et al. (2021). It is evident that the change in the number of new patents per period also shows a lumpy pattern. This finding further validates the relevance of R&D investment along the extensive margin.

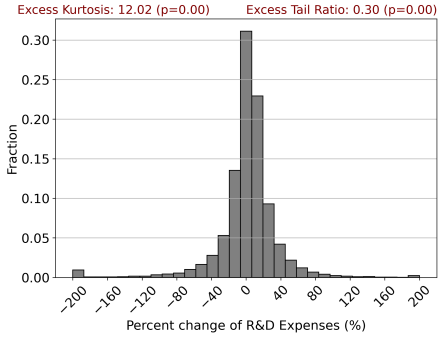
Figure B.1: Fact 1 Robustness



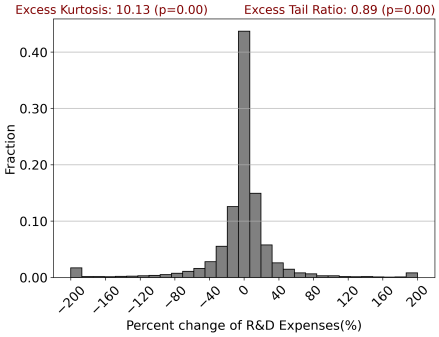
(a) Changes in the Number of New R&D Patents



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



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

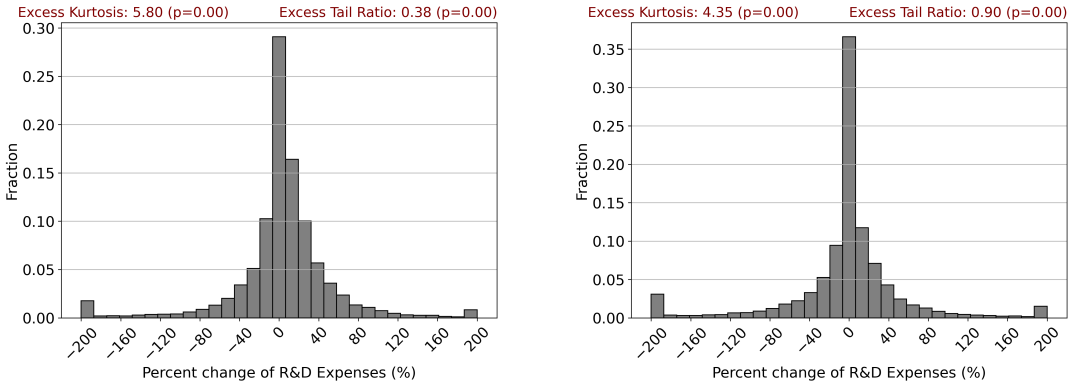


(d) Firms with Stable Scale past 3 Years

Note: 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.

Is Lumpy Innovation driven by Lumpy Changes in Firms' Scale? Another concern related to the lumpy pattern observed in Fact 1 is whether the observed lumpiness in R&D expenditure is merely a reflection of volatile changes in firms' scale (e.g., firms' sales). Firstly, sales may also manifest the pattern of spike changes. Should firms allocate R&D investment proportional to sales, this approach could inherently induce lumpy changes in R&D expenditures. Secondly, high sales might prompt firms to increase R&D spending from an accounting perspective to capitalize on tax incentives, thus this might contribute to the observed lumpy changes in R&D expenditures.

Figure B.2: Fact 1 Robustness: Financial Friction



(a) Low Leverage (b) High Leverage

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

To address this concern, Figure B.1b presents the distribution of the percentage change in the R&D-to-sales ratio. This approach normalizes R&D spending by firm sales, thus controlling for volatility in R&D inherited from sales. The main finding of lumpy R&D remains unchanged: The distribution of the change in the R&D-to-sales ratio still exhibits patterns of fat tails and a high concentration of values around zero.

Figure B.1d further examines whether the observed lumpy pattern is driven by a mismatch in the timing of lumpy changes in R&D and lumpy changes in firm scale. To address this, we consider a subsample of firms with relatively stable scale, defined as those whose sales have changed by less than 10% over the past three years. Figure B.1d shows that the lumpy pattern persists within this stable-scale subsample, suggesting that our results are not simply a consequence of lumpy changes in firm size.

Is Lumpy Innovation driven by Financial Friction? A potential concern regarding the findings of Fact 1 is that the lumpy pattern observed in R&D expenditures might be attributable to financial frictions rather than purely to fixed costs. We contend that this is not supported by the data. To investigate this issue, we analyze the changes in R&D expenditure across firms with differing levels of financial leverage, specifically comparing those with low leverage to those with high leverage, as measured by the debt-to-equity ratio. If financial frictions were indeed driving the lumpiness, we would expect firms with higher leverage, which face greater financial constraints, to exhibit more pronounced lumpy be-

Table B.2: Correlation of R&D Spike Definitions with GDP Growth

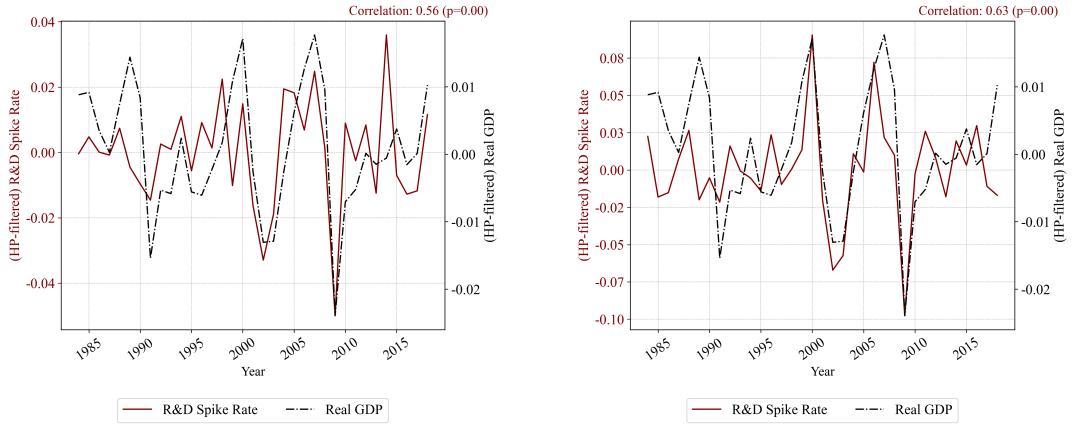
Spike Definition	Corr. GDP Growth	Corr. HP-filtered GDP	Average Spike Rate
$\Delta R\&D > 0.2$	0.45	0.37	0.09
$\Delta R\&D > 0.4$	0.54	0.40	0.04
$\Delta R\&D > 0.6$	0.54	0.36	0.02
$\Delta R\&D > 0.8$	0.50	0.33	0.01
$\Delta R\&D > 1.0$	0.35	0.26	0.01

havior. However, as depicted in Figure B.2, both low and high leverage firms display a similar degree of lumpiness in their R&D expenditure patterns. This similarity indicates that financial leverage is unlikely to be the primary factor behind the lumpiness, thereby reinforcing the argument that fixed costs play a more crucial role.

Procyclicality Given Alternative Definitions of Spikes In the main text, we identify an R&D observation as a spike if the R&D expense increases by at least 40%. Table B.2 reports several alternatives. The results suggest that raising the threshold for identifying spikes does not eliminate the pro-cyclicality of the spike rate.

The Lumpiness by Firms' Size of R&D A concern regarding the lumpiness of R&D is that it may only be observable in firms with smaller R&D expenditures. Firstly, methodologically, any change from zero to a positive R&D expense is identified as a spike. Secondly, firms with initially smaller R&D scales may exhibit more lumpiness in their R&D investments than larger firms, possibly due to the larger firms' experience and existing R&D resources. Figures B.1c and B.3 test this hypothesis. We define 'small R&D firms' as those whose average R&D size falls within the first and second quintiles, while 'large R&D firms' are categorized as those in the fourth and fifth quintiles. Figure B.1c demonstrates that large R&D firms also display significant lumpiness in their R&D investments. Figure B.3 compares the cyclicality of the spike rate between large and small R&D firms, with results indicating that the pro-cyclicality of the spike rate is similar for both large and small firms.

The Correlation with GDP Net of Aggregate R&D Spending Another concern related to the cyclicality of R&D investment is its inclusion as part of GDP. To verify that our findings on the cyclicality of lumpy R&D investment are not misleadingly influenced by its correlation with the R&D component of GDP, we replicate our analysis using GDP net of R&D spending. The same approach is also applied to unemployment data. Figure



(a) Small Firms

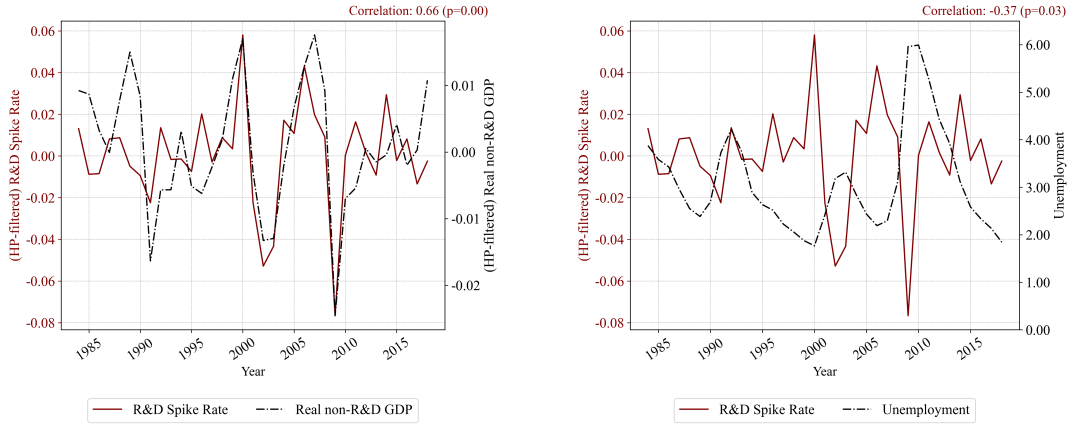
(b) Large Firms

Figure B.3: Fact 2 Robustness: Small v.s. Large R&D Firms. Note: 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.

B.4 shows that both analyses confirm the pro-cyclical pattern of R&D spikes, indicating that our results are robust and not affected by the inclusion of R&D spending in GDP calculations.

Without Entry/Exit One concern pertains to the interpretation of lumpy R&D activities based on accounting data from Compustat. Specifically, changes in reported R&D expenses from zero to a positive value may reflect changes in accounting policies rather than actual R&D activities. Regarding Fact 1, the calculation method for changes in R&D expenditure ensures that zero-to-positive changes are bounded by 200%. Even after excluding these bounded data points, the fat-tail nature of the histogram remains significant. Regarding Fact 2, Figure B.5a illustrates the relationship between the spike rate and GDP growth rate. Importantly, the spike rate excludes any zero-to-positive changes in R&D expenditure. As shown in the figure, the procyclicality of the spike rate remains significant, indicating that excluding these data points does not materially affect the evidence supporting Fact 2. This demonstrates the robustness of our findings.

Inaction The lumpy pattern highlighted in Fact 1 reveals a notable portion of firms with no changes in their R&D investment, indicating inaction. Following the definition by Cooper and Haltiwanger (2006), we represent this phenomenon by the inactive rate, which is defined as the fraction of firms with zero change in R&D expenses. As shown in Figure B.5b, the inactive rate exhibits significant counter-cyclicality, further supporting



(a) GDP net of R&D Spending

(b) Unemployment

Figure B.4: Fact 2 Robustness: pro-cyclicality Net of Aggregate R&D Spending. Notes: The GDP net of R&D Spending is calculated by substracting aggregate R&D spending from the real GDP. The series for the real GDP, aggregate R&D spending, and unemployment is from FRED.

the evidence of cyclicality in the lumpiness of R&D investment.

B.3 Data Appendix for the Great Twins

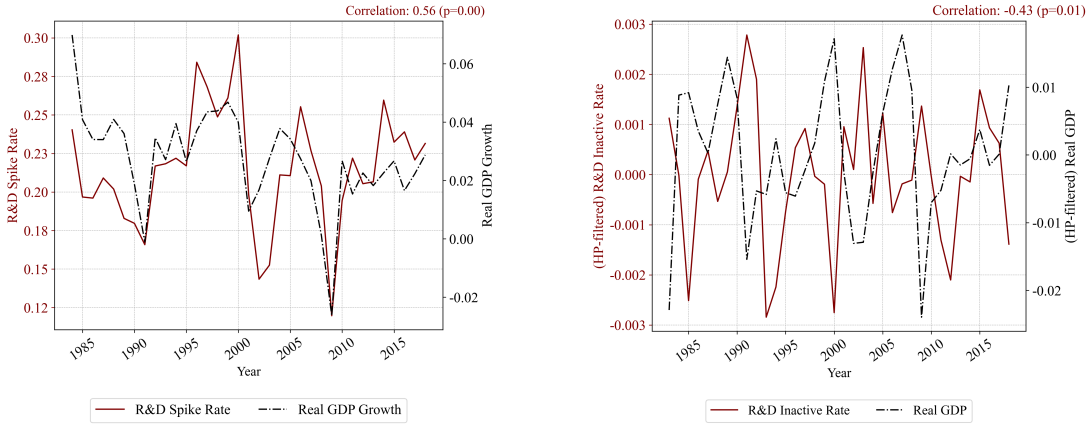
Measuring the Depth of Recession We formally define the gap between the actual GDP and the pre-recession trend level as follows:

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

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 the actual GDP observed at time t_1.

We define the depth of a recession following as the average deviation of actual GDP from its pre-recession trajectory:

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



(a) Without Entry

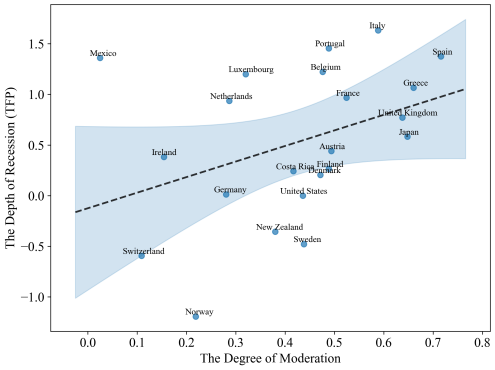
(b) Inactive Rate

Figure B.5: Fact 2 Robustness. Notes: The left figure shows the relationship between spike rate (excluding zero-to-positive changes) and GDP growth. The right figure plots the cyclicality of the inactive rate. The inactive rate is defined as the fraction of firms that do not change R&D expenses.

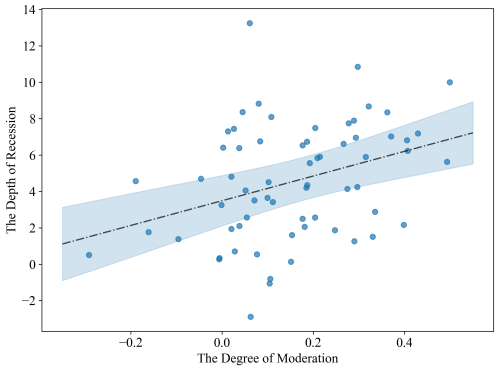
In practice, we select $K = 5$ because the recovery from the recession typically takes several years, during which the economy gradually aligns with its new trend. By taking the average, we further mitigate the fluctuations brought about by transitory shocks. Our empirical findings are robust to alternative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

The Great Twins Robustness Checks: TFP In the "Empirical Validation of the Model Predictions" section of the main text, we discovered that in cross-country data, the Great Twins phenomenon is evident—countries with lower volatility before 2008 experienced greater GDP losses due to the 2008 Great Recession. In our model, we attribute mid-term GDP losses to TFP losses. Figure B.6a examines the relationship between pre-2008 volatility and TFP losses caused by the 2008 Great Recession using the same cross-country sample. Consistent with our model prediction, we found that countries with lower pre-2008 volatility also experienced greater TFP losses due to the 2008 Great Recession.

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.6b 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.6b demonstrates that the Great Twins phenomenon exists across different countries and his-



(a) Great Twins: TFP



(b) Great Twins: Beyond 2008

Figure B.6: 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 ($p = 0.00$).

torical recessions. The positive correlation between pre-recession moderation and the loss of GDP remains significant, even when excluding the 2008 recession.

The Great Twins Robustness Checks: Sample Selection In the main text, we highlight that the empirical relationship of the Great Twins is significant in countries where R&D expenditure exceeds 1% of GDP, while the same relationship is not evident in countries with R&D expenditure below 1% of GDP. To ensure the robustness of this finding, we conduct additional checks on the sample selection where R&D expenditure exceeds 1% of GDP. As shown in columns 1 to 5 of Table B.3, we observe that this relationship is not apparent in the full sample. However, as we progressively restrict the sample to countries with higher R&D-to-GDP ratios, the relationship becomes increasingly significant.

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

Table B.3: 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***	15.47***	6.73***	-1.08	12.70***	12.33**	13.09***	13.07***
	(2.43)	(5.25)	(2.43)	(2.83)	(3.26)	(4.41)	(3.94)	(4.37)
$\frac{(\text{Private Credit/GDP})_{2007}}{(\text{Private Credit/GDP})_{2000}}$							3.99	3.97
							(2.33)	(2.56)
Real interest rate$_{2007}$						0.35		0.12
						(0.43)		(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**	0.17	2.86**	4.94***	0.81	0.12	-4.23	-4.48
	(1.10)	(2.41)	(1.06)	(1.23)	(1.48)	(2.14)	(3.59)	(3.98)
N	26	18	36	39	24	21	21	19

Note: 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.

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. The results are presented in Table B.3 (column AC1-AC4). We find that including these two control variables does not alter the significance of the coefficient for moderation.

C Model Appendix

C.1 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 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.

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 the elasticity parameter $\nu = 1$.

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 ϱ_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.2 Equilibrium, Linearization, and the Solution Method

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.

Linearization Based on the discussion in the Online Appendix C.1, a significant advantage of the model presented in this paper is that we can solve it through linearization. In this section, we will outline the process of linearizing the model and provide detailed explanations on the solution methods for the NK and Production Block, as well as the Innovation Block.

The Linearized Equations in the NK and Production Block In this notation, lowercase letters are employed to denote the logarithmic counterparts of variables denoted by uppercase letters. For instance, $\tilde{c}_t$ represents $\log \tilde{C}_t$. We further define $G_t = \frac{A_t}{A_{t-1}}$.

Household's supply of labor:

$$\tilde{c}_t + \varphi l_t = \tilde{w}_t, \quad \tilde{c}_t + \eta n_t = \tilde{w}_t^{RD} \tag{C.1}$$

Household's consumption Euler equation:

$$\tilde{c}_t = E_t \tilde{c}_{t+1} - (i_t - E_t \pi_{t+1} - E_t g_{t+1} + E_t \Delta \varrho_{t+1} - \rho) \tag{C.2}$$

The Intermediate goods producer's demand for M_t:

$$log(\gamma) + (1+\gamma) l_t - (1+\gamma) \tilde{m}_t = \log \vartheta \tag{C.3}$$

The Intermediate goods producer's demand for L_t:

$$log(1+\gamma) - \gamma l_t + \gamma \tilde{m}_t = \tilde{w}_t - \tilde{p}_{t,m} \tag{C.4}$$

Retailer's price-setting problem (NKPC):

$$\log(\Pi_t - 1) = \beta E_t \log(\Pi_t - 1) + \frac{\mu}{\kappa(1-\mu)} \log(\tilde{P}_{t,m} - \frac{1}{\mu}) \tag{C.5}$$

Monetary policy:

$$i_t = \rho + g_{t+1} + \varphi_\pi \pi_t \tag{C.6}$$

The market clearing for the labor market is straightforward. Regarding goods market clearing, note that the Cobb-Douglas production technology implies $\tilde{Y}_{t,m} = \gamma \tilde{M}_t$. Therefore, the market clearing writes $\tilde{C}_t = (1-\gamma)\tilde{Y}_{t,m}$ and the linearized version writes:

$$\tilde{c}_t = \log(1-\gamma) + \tilde{y}_{t,m} \tag{C.7}$$

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.

Solution Method: Innovation Block For the Innovation Block characterized by the Bellman equations, we employ the Value Function Iteration method for resolution. This involves incorporating the equilibrium wage of skilled labor, w_t^{RD}, and the growth-adjusted stochastic discount factor, $\tilde{\Lambda}_{t,t+1}$, determined within the NK and Production Block. To preserve the integrity of the linear approximation schema, we address the Bellman equations by employing a linear solution approach (deterministic equivalence), with a first-order approximation around the steady state.

C.3 A Formal Representation of Compressing and Stretching

Lemma A.1. *Consider a shock shifts the adjustment threshold $\bar{s}_t$ left- or rightward by Δ units from the steady state value $\bar{s}^*$. Let $f_{-\Delta}$ and $f_{+\Delta}$ denote the density function of the distribution of $s_{i,t+1}$ given the $-\Delta$ and $+\Delta$ shock, respectively, and $\bar{s}_{t+1}$ denotes the adjustment threshold at*

$t+1$:

(i) ***(Compressing Effect of Positive Shocks)*** *In response to a $-\Delta$ shock that shifts $\bar{s}^*$ leftward, there exists $s^{-\Delta} > 0$ such that $\forall \bar{s}_{t+1} > s^{-\Delta}$:*

$$f_{-\Delta}(\bar{s}_{t+1}) < f^*(\bar{s}_{t+1}), \tag{C.8}$$

(ii) ***(Stretching Effect of Negative Shocks)*** *if the shock shifts $\bar{s}^*$ rightward, there exists $s^{+\Delta} > 0$ such that for $\bar{s}_{t+1} > s^{+\Delta}$:*

$$f_{+\Delta}(\bar{s}_{t+1}) > f^*(\bar{s}_{t+1}). \tag{C.9}$$

Proof. See Online Appendix D. □

C.4 The Lumpiness and Asymmetry

In the following, we analyze two specific distributions of idea inflow, $\psi(\cdot)$, where the analytical solution for the steady state distribution of $s_{i,t}$, f, can be solved analytically given $\phi = 1$. In both cases, we show that the degree of asymmetry is growing with the lumpiness. Lastly, we provide numerical results showing that this pattern exists for various distributions of idea inflow and values of ϕ. We define the degree of asymmetry by:

$$\text{ASY}_\varepsilon(\bar{s}) \equiv \frac{\int_{\bar{s}(1-\varepsilon)}^{\bar{s}} f(x)dx - \int_{\bar{s}}^{\bar{s}(1+\varepsilon)} f(x)dx}{\int_{\bar{s}(1-\varepsilon)}^{\bar{s}(1+\varepsilon)} f(x)dx} \tag{C.10}$$

Exponential Distribution Assume $\psi(s) = \frac{1}{\lambda}\exp\{-\lambda s\}$. To simplify notation, we set $\lambda = 1$, while the proof for other values of λ is similar.

Using equation (D.4), the steady state distribution of $s_{i,t}$ solves:

$$f(s) = \tilde{P}\exp\{-s\} + \int_0^s f(s)\exp\{-s+x\}dx \tag{C.11}$$

for $s \leq \bar{s}$. One can show $f(s) = \tilde{P}$ solves the equation above. Therefore, for $s > \bar{s}$:

$$f(s) = \tilde{P}\exp\{-s\} + \int_0^{\bar{s}} \tilde{P}\exp\{-s+x\}dx = \tilde{P}\exp\{\bar{s}\}\exp\{-s\}$$

Finally, recall $\tilde{P} = \int_{\bar{s}}^{+\infty} f(x)dx$, which suggests $\int_0^{\bar{s}} f(s)dx + \tilde{P} = 1$. Because $\int_0^{\bar{s}} f(s)dx = \tilde{P}\bar{s}$:

$$\tilde{P} = \frac{1}{\bar{s}+1} \tag{C.12}$$

By some algebra, we can show that the degree of asymmetry in this case can be written as:

$$\text{ASY}_\varepsilon(\bar{s}) = \frac{\bar{s} - 1 + \exp\{-\bar{s}\varepsilon\}}{\bar{s} + 1 - \exp\{-\bar{s}\varepsilon\}} \tag{C.13}$$

which is increasing in $\bar{s}$.

Hybrid Gamma Distribution Next, we allow $\eta_{i,t}$ to follow a hybrid gamma distribution, which is described as:

$$\psi(\eta) = \frac{1}{2}(1+x)\exp(-x) \tag{C.14}$$

Using equation (D.4), the steady state distribution of $s_{i,t}$ solves:

$$f(s) = \frac{1}{2}\tilde{P}(1+s)\exp(-s) + \int_0^s \frac{1}{2}f(s)(1+s-x)\exp\{-s+x\}dx \tag{C.15}$$

One can show that the solution:

$$f(x) = \begin{cases} \tilde{P}\left(\frac{4}{3} - \frac{1}{6}\exp\left(-\frac{3}{2}x\right)\right) & \text{for } s \leq \bar{s}, \\ \tilde{P}\left(\frac{4}{3}\exp(-\bar{s}) - \frac{1}{6}\exp\left(-\frac{1}{2}\bar{s}\right)\right)\exp(-s) & \text{for } s > \bar{s}. \end{cases} \tag{C.16}$$

Then it follows:

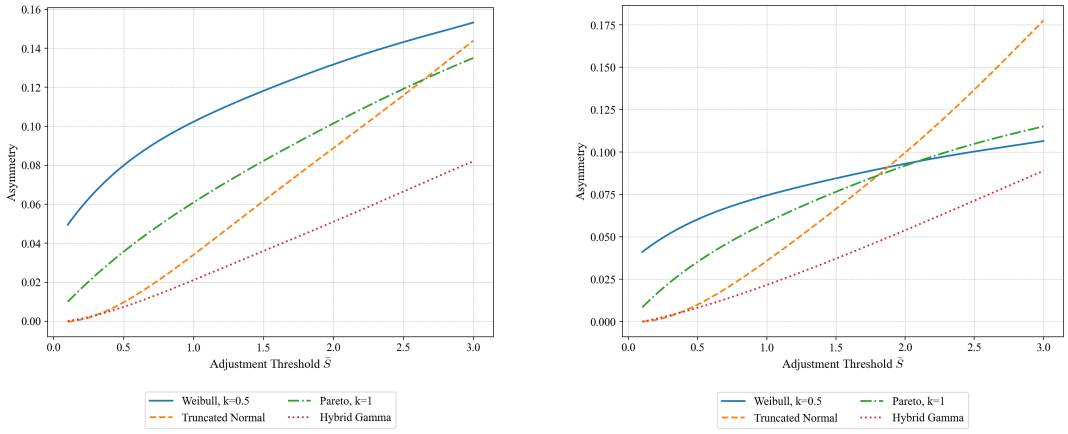
$$\int_{\bar{s}}^{\bar{s}(1+\varepsilon)} f(x)dx = \tilde{P}\Big[\frac{4}{3}(1-\exp(-\bar{s}\varepsilon)) - \frac{1}{6}\exp(-\frac{3}{2}\bar{s})\big(1-\exp(-\bar{s}\varepsilon)\big)\Big] \tag{C.17}$$

$$\int_{\bar{s}(1-\varepsilon)}^{\bar{s}} f(x)dx = \tilde{P}\Big[\frac{4}{3}\bar{s}\varepsilon - \frac{1}{6}\exp(-\frac{3}{2}\bar{s})\big(\frac{2}{3}\exp(\frac{3}{2}\bar{s}\varepsilon) - \frac{2}{3}\big)\Big] \tag{C.18}$$

One can show that for an approximation to the second order:

$$\text{ASY}_\varepsilon(\bar{s}) = \frac{\frac{2}{3} - \frac{7}{12}\exp(-\frac{3}{2}\bar{s})}{\frac{8}{3} - \frac{1}{3}\exp(-\frac{3}{2}\bar{s})} \cdot (\bar{s}\varepsilon) \tag{C.19}$$

Figure C.1: Asymmetry and Lumpiness



(a) $\phi = 0.9$ (b) $\phi = 0.8$

Note: We put $\varepsilon = 0.1$. The density function of the Pareto distribution is $\psi(x;k) = \frac{k}{x^{k+1}}$. The density function of the Weibull distribution is $\psi(x;k) = k\,(x)^{k-1}\,e^{-x^k}$.

which is increasing in $\bar{s}$.

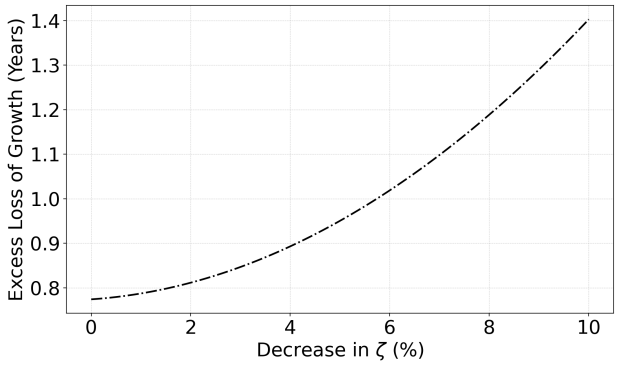
Numerical Analysis In most cases, deriving an analytical solution for the steady-state distribution function, f, is impossible. To verify if the relationship between increasing asymmetry and lumpiness generally holds, we conduct numerical exercises. These explore the association between asymmetry and lumpiness across various distributions of idea inflow and idea obsolescence rates. Figure C.1 demonstrates that the positive relationship between asymmetry and lumpiness persists across different model specifications.

C.5 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—defined as the number of years of trend growth lost due to the 2008–2009 shocks, conditional on the 2008 macroeconomic state, relative to the average

Figure C.2: The Great Twins when Ideas are Harder to Find



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—indicating a more pronounced and persistent scarring effect when ideas become harder to find.

C.6 Great Moderation in Lumpy Capital Investment

We first follow Baley and Blanco (2021) to 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:

$$y_{i,t} = (Z_t e_{i,t})^{1-\alpha} k_{i,t}^{\alpha}, \quad \text{where} \quad \alpha < 1. \tag{C.20}$$

Productivity components evolve according to:

$$z_{t+1} = \varepsilon_{t+1}^{z}, \quad \log e_{i,t+1} = \log e_{i,t} + \varepsilon_{i,t+1}^{e}. \tag{C.21}$$

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:

$$\theta_{i,t} = Z_t e_{i,t} \Theta(i_t, \Delta N_t). \tag{C.22}$$

where N_t is a Poisson counter with arrival rate λ:

$$\Theta(i_t, \Delta N_t) = \begin{cases} 0 & \text{if } i_t = 0 \text{ or } \Delta N_t = 1, \\ \theta^- & \text{if } i_t > 0 \text{ and } \Delta N_t = 0, \\ \theta^+ & \text{if } i_t < 0 \text{ and } \Delta N_t = 0. \end{cases} \tag{C.23}$$

Assuming no capital depreciation, capital stock evolves according to:

$$k_{i,t} = k_{i,t-1} + i_{i,t}. \tag{C.24}$$

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

$$V(k_{i,t-1}; Z_t, e_{i,t}) = \max_{i_{i,t}} \left\{ (Z_t e_{i,t})^{1-\alpha} k_{i,t}^{\alpha} - r k_{i,t} - \theta_{i,t} + \beta E_t[V(k_{i,t}; Z_{t+1}, e_{i,t+1})] \right\}. \tag{C.25}$$

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

$$i_{i,t} = k^*(Z_t, e_t) - k_{i,t-1}. \tag{C.26}$$

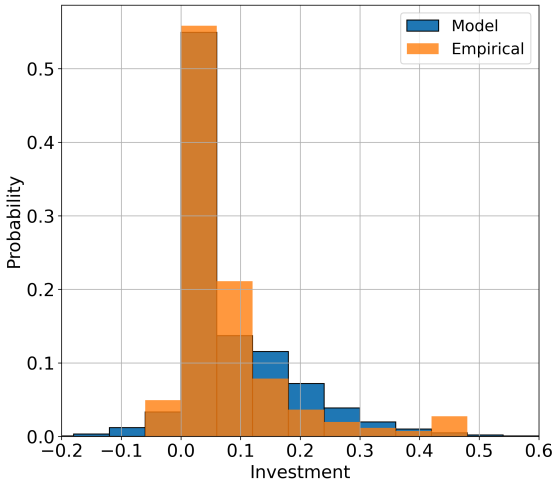
Following Baley and Blanco (2021), the optimal policy is determined by $\{k^-, k^*, k^+\} = e \times \{Zx^-, Zx^*, Zx^+\}$. We define the capital gap as $\tilde{x}_{i,t} = \log(x^*) - \log(\frac{k_{i,t}}{e_{i,t}})$.

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:

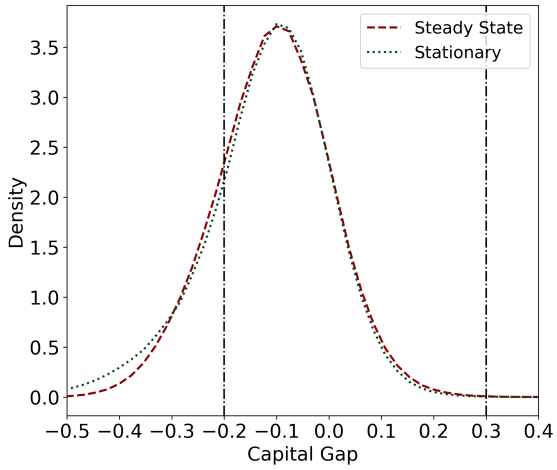
$$\tilde{x}_{i,t} = (\tilde{x}_{i,t-1} + \varepsilon^e_{i,t}) \cdot I\Big\{ \log(x^*) - \log(x^+) - z_t \leq \tilde{x}_{i,t-1} + \varepsilon^e_{i,t} \leq \log(x^*) - \log(x^-) - z_t \Big\}$$

Based on this parsimonious *Ss* model, Figure C.4 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 *Ss* framework.

Figure C.3: Calibrated Lumpy Investment Model



(a) Investment Distribution



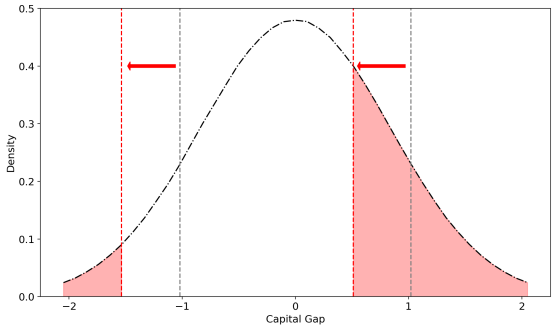
(b) Distribution of Capital Gap

Note: 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 asymetric: $\log(x^+) - \log(x^*) = 0.3$, $log(x^-) - \log(x^*) = -0.2$. The standard deviation of idiosyncratic productivity $\sigma^e = 0.26$. Figure C.3a 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.3b compares the steady state and stationary distribution of capital gap.

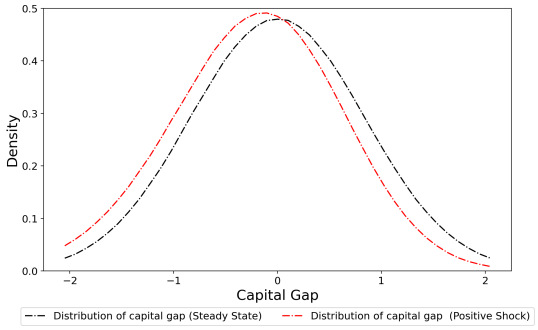
In a more realistic scenario, factors such as capital depreciation, partial irreversibility, and random opportunities for free adjustments may also influence the impact of aggregate shocks on the distribution of the capital gap. Therefore, in Figure C.3, we calibrate the parsimonious lumpy investment model to Compustat capital investment data. A comparison of the average distribution of the capital gap with the steady-state distribution of the capital gap in Figure C.3b 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.

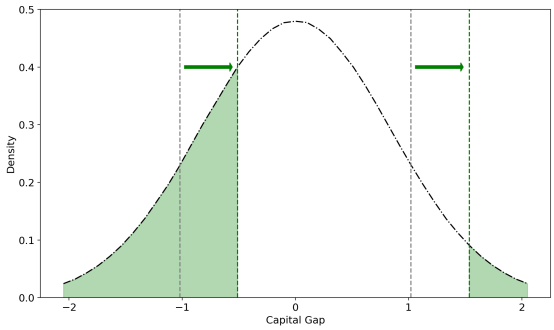
Figure C.4: The Effect of Aggregate Shocks on the Distribution of Capital Gap



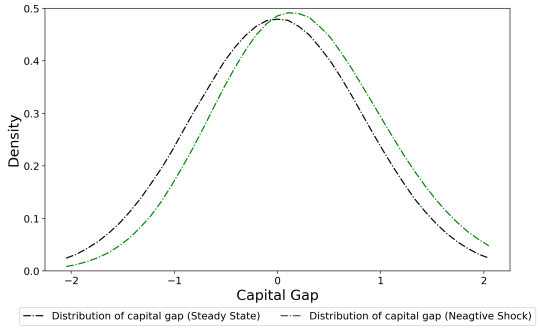
(a) Positive shocks shift inactive region to the left



(b) The distribution after a positive shock



(c) Negative shocks shift inactive region to the right



(d) The distribution after a negative shock

D Proof

D.1 The Proof of Proposition 1

Proof. We prove the proposition by first assuming there is $\bar{s}$, such that $\beta E[S(\phi\bar{s}+\eta)] = -c+\bar{s}J+\beta E[S(\eta)]$, and

$$S(s) = I\{s<\bar{s}\}(\beta E[S(\phi s+\eta)]) + (1-I\{s<\bar{s}\})(-c+sJ^A+\beta E[S(\eta)]) \tag{D.1}$$

Next we show that equation (D.1) is preserved by the recursive representation.

Denoting

$$f(s) = -c+sJ^A+\beta E[S(\eta)], \quad g(s) = \beta E[S(\phi s+\eta)] \tag{D.2}$$

First note that $f'(s) = J^A$ for all $s>0$.

Next, we rewrite $g(s)$ as

$$g(s) = \sum_k p_{1,k}(s)g_{1,k}(s)$$

where

$$p_{1,k}(s) = Pr(\phi s+\eta \le \bar{s}, \phi^2 s+\phi\eta+\eta_1 \le \bar{s}, ..., \phi^k s+\sum_i^k \phi^{i-1}\eta > \bar{s})$$

$$g_{1,k}(s) = \beta^k[-c+(\phi^k s+E[\sum_i^k \phi^{i-1}\eta | \sum_i^k \phi^{i-1}\eta > \bar{s}]) + \beta E[S(\eta)]]$$

therefore,

$$g'(s) = \sum_k p'_{1,k}(s)g_{1,k}(s) + \sum_k p_{1,k}(s)g'_{1,k}(s)$$

Note that $g'(s)$ is increasing and

$$\lim_{s\to\infty} g'(s) = \beta\phi J^A$$

which suggests $g'(s) < f'(s)$ for all $s \ge 0$. Hence, it preserve the property that for $s>\bar{s}$, $g(s)<f(s)$; and for $s<\bar{s}$, $g(s)>f(s)$, which completes the proof.

D.2 The Proof of Lemma 1

Proof. For the adjustment threshold $\bar{s} > 0$, the law for motion for $s_{i,t+1}$ is:

$$s_{i,t+1} = \phi s_{i,t} I\{s_{i,t} < \bar{s}\} + \eta_{i,t+1}. \tag{D.3}$$

We prove the case where $\phi = 1$ and the proof when $\phi < 1$ is similar. Denote ψ the density function of $\eta_{i,t}$, the implied equation for the steady-state distribution of $s_{i,t}$ is given by:

$$f^*(s) = \psi(s) \int_{\bar{s}^*}^{\infty} f^*(x)\, dx + \int_0^{\bar{s}^*} f^*(x)\psi(s-x)\, dx \tag{D.4}$$

To show that the absolute effect of compressing exceeds the effect of stretching, it is equivalent to show that the threshold $\bar{s}^*$ lies on the declining region of the steady state distribution of $s_{i,t}$.

Taking the first derivative on equation (D.4) around the thrshold $\bar{s}^*$:

$$\left.\frac{df^*(s)}{ds}\right|_{s=\bar{s}^*} = \psi'(\bar{s}^*) \int_{\bar{s}^*}^{\infty} f^*(x)\, dx + \int_0^{\bar{s}^*} f^*(x)\psi'(\bar{s}^*-x)\, dx \tag{D.5}$$

Because the idea inflow is non-negative and ψ is a decreasing function, it follows that:

$$\left.\frac{df^*(s)}{ds}\right|_{s=\bar{s}^*} < 0. \tag{D.6}$$

□

D.3 The Proof of Lemma A.1

Proof. The density after a $-\Delta$ shock to the adjustment threshold is implied by following transition equation:

$$f_{-\Delta}(s) = \psi(s) \int_{\bar{s}^*-\Delta}^{\infty} f^*(x)\, dx + \int_0^{\bar{s}^*-\Delta} f^*(x)\psi(s-x)\, dx. \tag{D.7}$$

And the density after a $+\Delta$ shock:

$$f_{+\Delta}(s) = \psi(s) \int_{\bar{s}^*+\Delta}^{\infty} f^*(x)\, dx + \int_0^{\bar{s}^*+\Delta} f^*(x)\psi(s-x)\, dx. \tag{D.8}$$

Note that it holds:

$$f_{-\Delta}(s) - f^*(s) = -\int_{\bar{s}^*-\Delta}^{\bar{s}^*} f^*(x)[\psi(s-x) - \psi(s)]\, dx < 0, \tag{D.9}$$

because $\psi(s)$ is strictly decreasing, ensuring $\psi(s-x) - \psi(s) > 0$.

Similarly,

$$f_{+\Delta}(s) - f^*(s) = -\int_{\bar{s}^*}^{\bar{s}^*+\Delta} f^*(x)[\psi(s) - \psi(s-x)]\, dx > 0, \tag{D.10}$$

because $\psi(s-x) - \psi(s) > 0$. □

D.4 The Proof of Proposition 2

Given $\varrho_t = (-1)^t\sigma$, the adjustment threshold $\bar{s}_t = \bar{s}^*(1 + (-1)^t\Delta(\sigma))$. We consider the transition from the steady state distribution to the stationary.

$$f^{(1)}_{-\sigma}(s) = \psi(s)\int_{\bar{s}^*(1-\sigma)}^{\infty} f^*(x)\, dx + \int_0^{\bar{s}^*(1-\sigma)} f^*(x)\psi(s-x)\, dx. \tag{D.11}$$

By equation (D.9), if ψ is monotonically decreasing, there is some $s^\circ < \bar{s}^*$ such that $f^{(1)}_{-\sigma}(s) < f^*(s)$ for $s > s^\circ_{(1),-\sigma}$. Now, consider

$$f^{(1)}_{+\sigma}(s) = \psi(s)\int_{\bar{s}^*(1+\sigma)}^{\infty} f^{(1)}_{-\sigma}(x)\, dx + \int_0^{\bar{s}^*(1+\sigma)} f^{(1)}_{-\sigma}(x)\psi(s-x)\, dx. \tag{D.12}$$

Because $f^{(1)}_{-\sigma}(s) < f^*(s)$ for $s > s^\circ$ and ψ is decreasing, it holds

$$\int_{\bar{s}^*(1+\sigma)}^{\infty} f^{(1)}_{-\sigma}(x)\, dx < \int_{\bar{s}^*(1+\sigma)}^{\infty} f^*(x)\, dx, \tag{D.13}$$

and

$$\int_0^{\bar{s}^*(1+\sigma)} f^{(1)}_{-\sigma}(x)\psi(s-x)\, dx < \int_0^{\bar{s}^*(1+\sigma)} f^*(x)\psi(s-x)\, dx, \tag{D.14}$$

for $s > s^\circ$. Equation (D.14) is because:

$$\int_0^{\bar{s}^*(1+\sigma)} f^{(1)}_{-\sigma}(x)\psi(s-x)\, dx = \int_0^{s} f^{(1)}_{-\sigma}(x)\psi(s-x)\, dx = \int_0^{s} f^{(1)}_{-\sigma}(s-x)\psi(x)\, dx \tag{D.15}$$

and

$$\int_0^s f_{-\sigma}^{(1)}(s-x)\psi(x)\,dx > \int_0^s f^*(s-x)\psi(x)\,dx \tag{D.16}$$

because the LHS integral gives more weight to $f_{-\sigma}^{(1)}(s-x)$ on the range where $f_{-\sigma}^{(1)}$ is larger than f^*. Therefore, by equation (D.13) and (D.14), it holds $f_{+\sigma}^{(1)}(s) < f^*(s)$ for $s > s_{(1),+\sigma}^{\circ}$, for some $s_{(1),+\sigma}^{\circ} < \bar{s}^*$.

Now consider round k, if it holds $f_{+\sigma}^{(k-1)}(s) < f^*(s)$ for $s > s_{(k-1),+\sigma}^{\circ}$,

$$f_{-\sigma}^{(k)}(s) = \psi(s)\int_{\bar{s}^*(1-\sigma)}^{\infty} f_{+\sigma}^{(k-1)}(x)\,dx + \int_0^{\bar{s}^*(1-\sigma)} f_{+\sigma}^{(k-1)}(x)\psi(s-x)\,dx. \tag{D.17}$$

it is easy to show that for some $s_{(k),-\sigma}^{\circ} > 0$, $f_{-\sigma}^{(k)}(s) < f_{+\sigma}^{(k-1)}(s)$ holds for $s > s_{(k),-\sigma}^{\circ}$. Similar to the logic of showing (D.13) and (D.14), it can be shown that for

$$f_{+\sigma}^{(k)}(s) = \psi(s)\int_{\bar{s}^*(1+\sigma)}^{\infty} f_{-\sigma}^{(k)}(x)\,dx + \int_0^{\bar{s}^*(1+\sigma)} f_{-\sigma}^{(k)}(x)\psi(s-x)\,dx, \tag{D.18}$$

it holds that $f_{+\sigma}^{(k)}(s) < f^*(s)$ for $s > s_{(k),+\sigma}^{\circ}$, for some $s_{(k),+\sigma}^{\circ} < \bar{s}^*$.

Therefore, by induction, for any k, it holds for $s > \max\{s_{(k),+\sigma}^{\circ}, s_{(k),+\sigma}^{\circ}\}$,

$$\frac{f_{-\sigma}^{(k)}(s) + f_{+\sigma}^{(k)}(s)}{2} < f^*(s). \tag{D.19}$$

Further, because $\bar{s}^* > \max\{s_{(k),+\sigma}^{\circ}, s_{(k),+\sigma}^{\circ}\}$:

$$\frac{f_{-\sigma}^{(k)}(\bar{s}^*) + f_{+\sigma}^{(k)}(\bar{s}^*)}{2} < f^*(\bar{s}^*), \tag{D.20}$$

which completes the proof.