Lumpy R&D and the Great Twins: Great Moderation and Great Recession*

Lixing Wang† Donghai Zhang‡

May 6, 2024

Abstract

We document that firm-level R&D investments are lumpy, characterized by episodes of inaction and large increases in R&D expenses. Furthermore, the share of firms making spike adjustments in their R&D expenses is pro-cyclical and explains 75% of the pro-cyclicality of aggregate R&D. Motivated by these empirical findings, we develop a medium-run business cycle model with heterogeneous innovative firms and fixed adjustment costs, featuring lumpy R&D. Within this framework, we identify a novel mechanism that gives rise to the 'Great Twins' hypothesis: the Great Moderation—characterized by reduced economic fluctuations until 2008—caused the observed severity of the subsequent Great Recession. We provide micro- and macro-level evidence supporting this hypothesis. Finally, we discuss the policy implications of our findings, including the novel tail-risk and variance tradeoff faced by central banks.

Keywords: Heterogeneous Innovators, Lumpy R&D, Distribution, Great Twins, Medium-run Business Cycles
JEL Classification: E32, O31, O33

*We thank Diego Comin, Ryan Chahrour, John Fernald, Pablo A. Guerrón-Quintana, Franck Portier, Yongseok Shin, Michael Song as well as the seminar participants at the University of Bonn, Singapore Management University, National University of Singapore, Fudan University, SUFE, Academia Sinica, and Virtual East Asian Macroeconomics Seminar, Southern Workshop in Macroeconomics for their suggestions and discussions. All remaining errors are our own.

†University of Bonn, lixing.wang@uni-bonn.de
‡National University of Singapore and University of Bonn, donghai.d.zhang@gmail.com

1 Introduction

Research and Development (R&D) projects create new knowledge and technology that promote economic growth. Business cycles affect the profitability of new technology, and, therefore, influence innovators' R&D decisions. This insight, coupled with the observed pro-cyclicality of *aggregate* R&D expenditure, has motivated a growing literature on medium-run business cycles (Comin and Gertler, 2006) that emphasizes the effects of business cycles on medium-run economic growth via their impact on R&D. 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' 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 R&D 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. 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 R&D for medium-run business cycles?

To formally motivate our theoretical framework, we document three novel facts. First, the distribution of changes in R&D expense at firm-level is heavily populated by zeros/small values and large values—lumpy R&D. Second, the fraction of firms making big adjustments (or spike adjustments) in R&D expense—the spike rate—is pro-cyclical. Third, the pro-cyclicality of the *spike* adjustment is the primary source of the pro-cyclicality of *aggregate* R&D expense documented in the existing literature.

Motivated by these findings, we develop a medium-run business cycle model with *heterogeneous* innovators who are subject to fixed adjustment costs in setting up new R&D projects, giving rise to lumpy R&D. The model can replicate the firm-level distribution and the cyclical behavior of spike adjustments observed in the data due to innovators' decisions along the extensive margin. The model offers a novel insight: the Great Twins hypothesis. Specifically, the Great Moderation—characterized by reduced economic fluctuations until 2008—caused the observed severity of the subsequent 2008 Great Recession.

Crucial to this finding is a novel mechanism (Proposition 3) in the model: the asymmetric effects of shocks on the distribution of innovators at the adjustment threshold. The Great Twins hypothesis reveals a new tail-risk and volatility tradeoff faced by central banks: while good monetary policy leads to a smaller unconditional volatility, it results in a larger drop in GDP conditional on a recession. Additionally, our model predicts a Recession trap, where nominal demand stimulus policy is less effective in a recession than in a boom. We provide policy recommendations to address the challenges posed by the Great Twins and Recession trap. Last but not least, we present macro-evidence that supports the Great Twins hypothesis in the data and micro-evidence that supports the mechanism behind the Great Twins.

In more detail, we develop a medium-run business cycle model with endogenous growth 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 an exogenous amount of raw ideas. Innovators are heterogeneous in terms of the amount of new ideas they receive. 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, the innovative 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 R&D decision. To convert R&D ideas into new varieties of intermediate input, 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 varieties 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 R&D decisions. Business cycle fluctuations affect the profits of new varieties, 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. Fur-

thermore, we fit the model to the evolution of real GDP to estimate business cycle shocks, demonstrating the model's capability to generate the Great Recession with shocks of reasonable size (within two standard deviations). Next, we turn to the main implications of the model.

The model offers a novel 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), 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 good policy that reduces volatility inadvertently has an unintended negative consequence of increased left tail risk.

Two highlighted features drive the result. The first is the fixed cost of setting up a new project, which gives rise to changes in R&D expenses along the extensive margin. Second, innovating firms are heterogeneous. As a result, a key state variable is the distribution of the stock of raw ideas across firms. Business cycles affect innovators' extensive margin decisions, and therefore, alter the distribution of the stock of raw ideas. This distribution, as a state variable, in turn, affects the propagation of shocks. The Great Moderation describes the moderated macroeconomic volatility in the U.S. economy before the 2008 Great Recession. The reduced macroeconomic volatility stretches out the distribution of total ideas across innovators, making the distribution *more fat-tailed*. As a result, innovators are more likely to adjust along the extensive margin, leading the economy to a greater chance of extreme adverse outcomes associated with declined R&D innovations. This moderation-induced fat-tailed distribution is due to a novel mechanism offered by our model: asymmetric effects of shocks on the distribution, which holds true qualitatively (see Proposition 3). To the best of our knowledge, such a mechanism has not been documented before.

Additionally, our model suggests that nominal demand stimulus policies become less effective during recessions–when they're most needed, a situation we call the 'Recession Trap.' Specifically, a cut in the nominal interest rate has only one-third of its boom-time effect in a recession.

We offer policy recommendations to address the Great Twins and the Recession Trap. We demonstrate that R&D subsidies can reduce both the volatility and left-tail risk of

the economy. Moreover, in a recession, if an interest rate cut is accompanied by an R&D subsidy policy, the net effects of the interest rate cut would be nearly doubled compared to the scenario without an R&D subsidy.

Last but not least, we provide macro- and micro-level evidence supporting the Great Twins hypothesis and validating its underlying mechanism. Relying on cross-country variations, we illustrate that a country experiences a deeper 2008 recession if it has undergone a more moderated economy in the preceding decade. This finding is not unique to the Great Recession: generalizing the empirical analysis to cover all recessions in the post-WW2 sample yields the same pattern. Moreover, 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.

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). 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. 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). Differing from these papers, we emphasize heterogeneity in the propensity to innovate, and our paper establishes implications beyond variables in the steady state. Specifically, we introduce a new model to understand how business cycle fluctuations affect the distribution of innovators' propensity to innovate and, importantly, how these distribution changes influence the propagation of shocks. While 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. We present a novel model that connects business cycle fluctuations with innovators' R&D decisions along the extensive margin.

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), and more recently, Koby and Wolf (2020), Winberry (2021), Baley and Blanco (2021), Gnewuch

and Zhang (2023), and Lippi and Oskolkov (2023). To the best of our knowledge, the new mechanism in our model is unique within the context of lumpy R&D and has not been previously documented in the literature on lumpy capital investment. Overall, our paper introduces a mechanism that leads to a novel insight–the 'Great Twins' hypothesis–which suggests a trade-off between tail-risk variance that policymakers should consider when formulating stabilization policies. This result builds upon the mechanism by which distribution affects the propagation of business cycle shocks, a feature extensively emphasized within the heterogeneous agents literature. This literature examines the distribution of heterogeneous households (Krusell et al. 2010, McKay and Reis 2016, Krueger et al. 2015, Auclert 2019, Broer et al. 2020, Gornemann et al. 2016, Kaplan et al. 2018, Bayer et al. 2019) and firms (Khan and Thomas 2008, Reiter et al. 2013, Sedláček and Sterk 2017, Ottonello and Winberry 2020, Jungherr et al. 2022, Cloyne et al. 2023, Lee 2023, Tang and Zhang 2022, 2023, and Gnewuch and Zhang 2023). In our setting, firms' endogenous decisions affect the distribution of unused raw ideas. In this context, the novelty of the Great Twins result lies in its emphasis on how the *history of shocks* (Great Moderation) affects the distribution and thereby increases the risk in the left tail (Great Recession).

The empirical section of our study is related to the literature examining the cyclicality of R&D. Fatas (2000), Wälde and Woitek (2004), and Barlevy (2007) document that the aggregate R&D spending is pro-cyclical, with this pro-cyclicality being robust across countries and for different 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 and demonstrating that the pro-cyclicality of aggregate R&D is primarily explained by this extensive margin.

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 non-linear mechanism in our model is an *endogenous* feature. Moreover, different from the existing literature, we emphasize R&D investment along the extensive margin, consistent with our empirical findings that the pro-cyclicality of R&D investments is primarily due to the extensive margin.

2 Empirical Evidence: Lumpy R&D and its Cyclicality

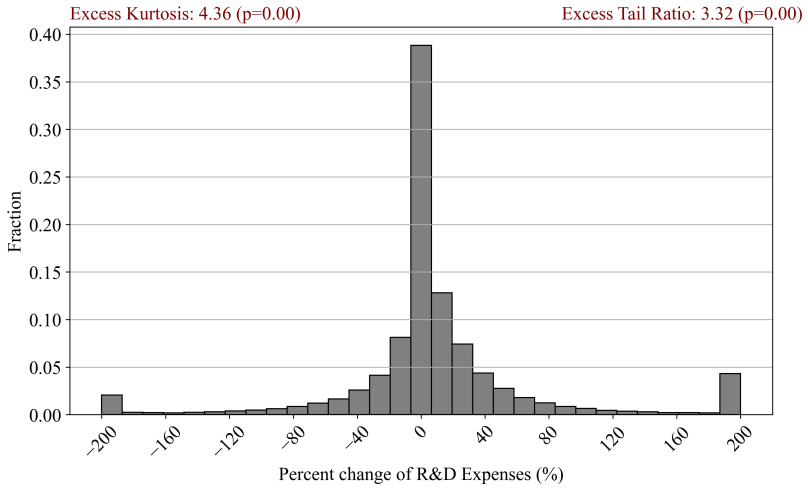
This section presents novel empirical evidence that motivates the model introduced in the subsequent section. Firstly, we document that the R&D investments of innovative firms exhibit a lumpy pattern characterized by a substantial proportion of spike adjustments and inaction. Additionally, the fraction of firms making significant adjustments to R&D spending–the spike rate–is pro-cyclical and plays a crucial role in explaining the cyclicality of aggregate R&D investment.

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 innovative firms. In the empirical analysis, a firm is classified as an innovative 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[1]. Our sample consists of 9,166 innovative firms, representing 42.8% of the companies in the Compustat universe. The panel is unbalanced, with each firm averaging 16.6 observations.

The primary variable in our study is Research and Development 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 and other costs associated with the devel-

[1]A concern regarding mergers and acquisitions is that a firm's acquisitions may incorporate the R&D departments of the acquired companies, leading to abrupt changes in the firm's R&D expenditures. To address this, we exclude observations involving acquisitions and focus on the firm's R&D behaviors.

Figure 1: Distribution of Percent Change of R&D Expenses



Note: The figure plots the firm-level distribution of the percent change in R&D expenses. To address the undefined change when firms start with zero R&D expenses, we calculate the percent change using the formula: percent change of R&D expenses $= (\text{R\&D expense}_t - \text{R\&D expense}_{t-1})/(\frac{1}{2}\text{R\&D expense}_t + \frac{1}{2}\text{R\&D expense}_{t-1})$. Excess kurtosis measures the deviation of the fourth standardized moment from a normal distribution. Similarly, the excess tail ratio compares the 80%-quantile-to-95%-quantile ratio to that of a normal distribution. The p-value is determined through bootstrapping.

opment of the concept, 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.

Formalizing the Concept of Lumpy R&D To formalize the concept of lumpy R&D, Figure 13 plots two examples of the evolution of R&D expenses for actual firms in the data. Most of the time, *changes* in R&D expenses are small. However, occasionally, innovative firms make a spike adjustment in R&D (indicated by the red dot). We refer to these observations as evidence of lumpy R&D.

Formally, lumpy R&D postulates that the distribution of *changes* in R&D ($\Delta \log RD_{i,t}$) is heavily populated by zeros/small values and large values. Is lumpy R&D a general phenomenon in the data? This question is addressed in the following section.

2.1 Fact 1: Lumpy R&D

Figure 1 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 R&D: R&D invest-

ment is characterized by no or small changes and large spikes. The excess kurtosis of this distribution is 4.36 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, representing the deviation in the distribution's 80th-quantile to 95th-quantile ratio from that of a normal distribution, is 3.32, markedly exceeding the null hypothesis value of 0, confirming the fat-tail pattern.

To further understand the concept of lumpy R&D, we establish a connection between the two cases of R&D expense evolution illustrated in Figure 13 and the distribution of the percent change in R&D expenses plotted in Figure 1. In Figure 13, the red dots highlight the period when firms make a spike adjustment in R&D expenses. These periods lie in the right tail of the distribution of the percent change of R&D expenses plotted in Figure 1. In periods following a spike adjustment, firms make relatively small changes in R&D expenses. These small changes are captured by the region surrounding zero in Figure 1.

In summary, the lumpiness of R&D that we emphasize stresses the significant presence of two key regions: the inaction region—characterized by zero or small changes in R&D expenses, and the spike region—representing spike adjustments.

2.2 The Cyclicality of Lumpy R&D

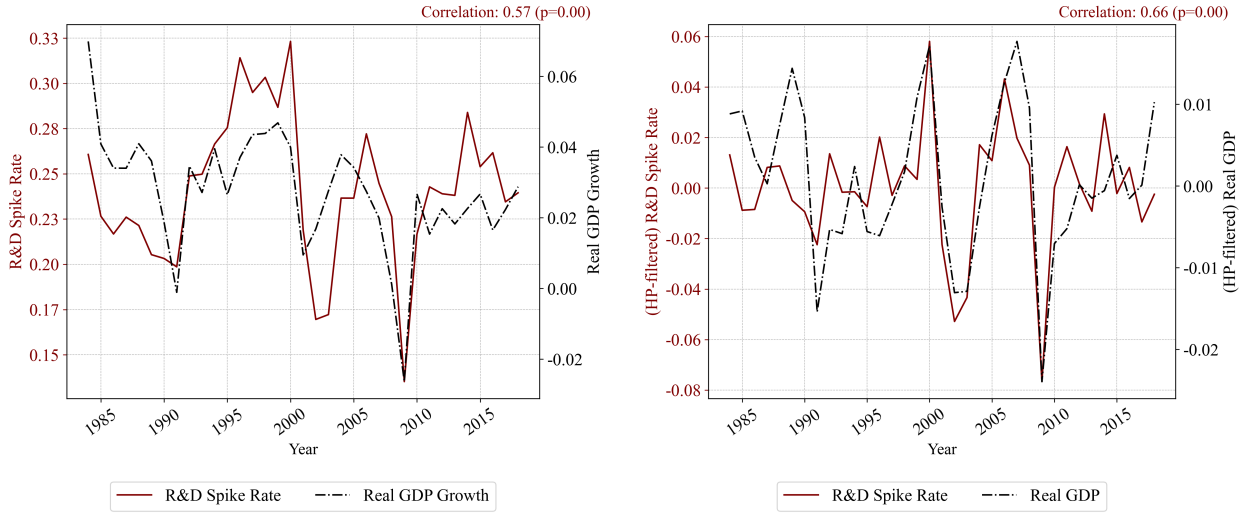
The existing literature has documented that aggregate R&D expenditure is pro-cyclical. Given the lumpy nature of changes in R&D expenses, there are two types: sizable ones (referred to as R&D spikes) and smaller ones. Two natural questions arise: Are R&D spikes cyclical? Which type of changes in R&D expenses drives the aggregate R&D cyclicality documented in the literature?

Fact 2: R&D spikes are pro-cyclical To address these questions, we first define R&D spikes as increases in R&D expenses by more than 20%. In the subsequent discussion, unless otherwise specified, R&D spikes specifically refer to positive spikes in R&D expenses.[2] The R&D spike *rate* is defined as the fraction of innovative firms (at time t) showing spikes in their R&D expenses. Our findings are robust to alternative choices of the threshold.

Figure 2a plots the evolution of the fraction of R&D spikes against real GDP growth

[2]We focus on positive R&D spikes because negative spikes can result from budget cuts or project shutdowns, often tied to the business cycle, or from completing R&D projects, which isn't necessarily cycle-related. Figure 3b shows that negative spikes contribute less to R&D's pro-cyclicality compared to positive spikes.

Figure 2: The Pro-cyclicality of R&D spikes



(a) R&D Spike Rate and Real GDP growth

(b) Filtered R&D Spike Rate and Real GDP

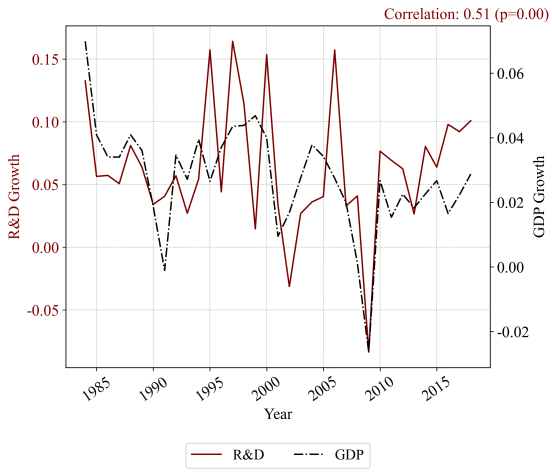
Note: The figures plot the time series of the R&D spike rate, defined as the fraction of R&D spikes, and real GDP from 1984 to 2018. A firm-year observation is identified as an R&D spike observation if the change in R&D expenses for that firm is greater than 20%. The spike rate is the fraction of firms that experienced an R&D spike in that period. In panel (b), both series are filtered by the HP filter, with $\lambda = 6.25$.

rate across time. The correlation between these two series is 0.57 (p-value = 0.00), suggesting that firms usually choose to make significant changes in R&D expenses during economic booms. This empirical pattern also exists if we filter both series using HP filters. Figure 2b shows that the two series align well with each other, with a correlation of 0.66 (p-value = 0.00).

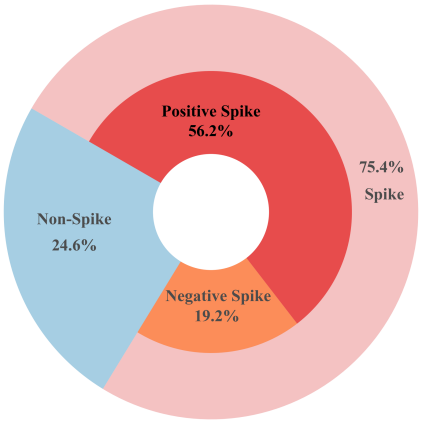
Figure 1 is characterized by a significant fraction of firms exhibiting inaction in R&D investment. Following Cooper and Haltiwanger (2006), we define the fraction of firms with zero change in R&D expenses as the inactive rate. Figure 15 demonstrates that the inactive rate is significantly counter-cyclical, providing further evidence of the cyclicality of lumpiness in R&D investment.

Fact 3: R&D Spikes Drive the pro-cyclicality of Aggregate R&D Previous studies (e.g., Barlevy 2007) have documented the pro-cyclicality of aggregate R&D. This pattern is also evident in our data. By aggregating R&D spending across firms, Figure 3a illustrates that aggregate R&D spending is strongly pro-cyclical, with a correlation of 0.51 (p-value

Figure 3: R&D Spikes Drive the pro-cyclicality of Aggregate R&D



(a) pro-cyclicality of Aggregate R&D



(b) Decompose the pro-cyclicality

Note: Figure 3a presents the time series for the growth of aggregate R&D expenditure alongside GDP growth. The aggregate R&D expenditure is calculated by summing up the R&D spending across firms in the compustat, while the GDP data is sourced from the FRED database. Figure 3b delves into the decomposition of the positive correlation observed in Figure 3a, categorizing the changes in aggregate R&D into those stemming from spikes and non-spikes. Specifically, spike changes are defined as the absolute values of R&D growth exceeding 20%. Furthermore, the changes attributed to spikes are subdivided into those resulting from positive and negative spikes.

$= 0.00)$[3].

How significant is the impact of the lumpy nature of R&D investment at the micro-level in explaining the cyclicality of aggregate R&D expenditures? To address this question, we first note that the correlation between the growth of R&D and GDP growth can be decomposed into the parts due to spike changes and non-spike changes, as shown in Appendix C.2. The formula is based on the simple observation that $Cov\left(\frac{\Delta R\&D_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right) = Cov\left(\frac{\Delta R\&D^{Spike}_{t+1}}{R\&D_t} + \frac{\Delta R\&D^{non-Spike}_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right) = Cov\left(\frac{\Delta R\&D^{Spike}_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right) + Cov\left(\frac{\Delta R\&D^{non-Spike}_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right)$, where $R\&D^{\text{non-Spike}}_{t+1}, R\&D^{\text{Spike}}_{t+1}$ represents the changes of R&D due to spikes and non-spikes, respectively.

Figure 3b presents the results of the decomposition analysis, illustrating that spikes in R&D expenditure contribute 75.4% to the pro-cyclicality of aggregate R&D investment. In contrast, non-spike variations account for only 24.6% of this cyclical behavior. We can further decompose the aggregate changes in R&D caused by spike changes into

[3]In Figure 16a, we also show that the HP-filtered aggregate R&D in compustat is highly pro-cyclical, with a correlation of 0.62 (p-value=0.00).

those driven by positive and negative spikes. Both the positive and negative spikes contribute to the pro-cyclicality of the aggregate changes. During booms (recessions), the frequency of positive spikes increases (decreases), leading to a rise (decline) in aggregate R&D changes. Conversely, during booms (recessions), the incidence of negative spikes decreases (increases), resulting in fewer (more) negative changes. The further decomposition of spike changes shows that their significant influence predominantly stems from positive spikes. These positive spikes alone are responsible for 56.2% of the aggregate R&D's pro-cyclicality.

2.3 Discussions and Robustness Checks

We conclude the empirical section by providing our interpretation of the observed facts and presenting additional robustness checks.

Interpretation We interpret 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. According to this interpretation, Facts 2 and 3 suggest that innovators are more likely to innovate along the extensive margin in a boom than in a recession, and the pro-cyclicality of innovators' extensive margin decisions drives the pro-cyclicality of aggregate R&D expenses.

Section 3 presents a novel model to formulate these arguments and study the implications. The key model assumption is motivated by the observation that a new R&D project has a unique feature, in requiring a substantial initial setup cost, such as purchasing new laboratory equipment and organizational restructuring, followed by a stream of stable variable costs, such as wage bills. Due to this feature, innovative firms might not set up new R&D projects or adjust R&D expenses frequently. In other words, R&D investment decisions are lumpy.

Lumpy Patents Lumpy R&D might imply that the outputs of R&D, such as patents, should also exhibit a lumpy pattern. Figure 14b displays the firm-level distribution of changes in the number of new patents, as recorded in the USPTO database. 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.

Tax Incentive One of the concerns regarding the result from Figure 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 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.

Is Lumpy R&D driven by Lumpy Sales? Another concern related to the lumpy pattern observed in Figure 1 is whether the observed lumpiness in R&D expenditure is merely a reflection of volatile changes in 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.

To address this concern, Figure 14a 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.

Additional Results In the Online Appendix C.3, we provide a detailed discussion on additional robustness checks for Facts 1-3. Firstly, we show that our results are robust across various alternative definitions of R&D spikes. Secondly, we establish that the lumpiness and cyclicality of R&D are consistent across firms of different R&D expenditure sizes, addressing concerns that our findings might be disproportionately influenced by smaller R&D firms. Thirdly, we confirm that the pro-cyclicality of lumpy R&D is not a spurious outcome due to the inclusion of an aggregate R&D component in real GDP.

3 Model

The framework builds upon a medium-run business cycle model with endogenous TFP growth (Comin and Gertler 2006 and Comin et al. 2009). In contrast to existing works

based on models with homogeneous innovative firms, we introduce *heterogeneous* innovative firms that need to pay a fixed cost for setting up a new R&D project. The inclusion of heterogeneous agents increases the model's complexity by incorporating the firm-level distribution of the stock of unadopted technologies, referred to as the idea stock distribution, as a state variable. This feature is crucial for the main predictions of the model.

Figure 17 provides an overview of our model system, broadly categorized into three main blocks: the business cycle block, the production block, and the innovation block. Within the business cycle block, households are modeled to allocate their income between consumption and savings, the latter in the form of government bonds. Consumption comprises a composite of consumption goods, assembled by retailers using intermediate goods. Moving to the production block, this features a representative producer of intermediate goods employing labor and a composite of intermediate inputs in the production process. The total number of varieties of intermediate inputs, denoted as A_t, increases over time (due to innovations), contributing to the endogenous growth of the economy. Each variety is produced by a distinct intermediate input producer, and the market for these intermediate inputs is designed to be monopolistically competitive, providing a monopolistic profit margin for each producer. In the innovation block, producers acquire the rights to manufacture specific varieties of intermediate input from their respective innovators. In return, these innovators are entitled to a portion of the production profits, representing the enactment of intellectual property rights.

We now discuss each component of the model in detail. Section 3.1 covers the production block. Section 3.2 introduces the innovation block, where our model's core mechanism is explained. Section 3.3 concludes the model description with the business cycle block. Section 3.4 presents the corresponding market clearing conditions that define the general equilibrium.

3.1 The Production Block

The production block refers to the production of intermediate goods $Y_{t,m}$ and the *production* of different varieties of intermediate input $M_t(j)$. The latter differs from the *creation* of new varieties of intermediate input, which will be discussed later in Section 3.2.

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

$M_t(i)$:

$$M_t = \left(\int_0^{A_t} M_t(j)^{\frac{1}{\vartheta}} dj \right)^{\vartheta}, \tag{3.1}$$

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., expansion in A_t, is the source of productivity growth. The intermediate inputs $M_t(i)$ are created by innovators, who will be introduced later in Section 3.2.

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 $M_{t,j}$:

$$M_{t,j} = \left(\frac{P_{t,j}^M}{P_t^M} \right)^{\frac{\vartheta}{1-\vartheta}} M_t, \tag{3.2}$$

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:

$$Y_{t,m} = L_t^{1-\gamma} M_t^{\gamma}, \tag{3.3}$$

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:

$$\max_{M_t, L_t} P_{t,m} Y_{t,m} - W_t L_t - P_t^M M_t, \tag{3.4}$$

subject to production function (3.3). $P_{t,m}$ is the price of the intermediate good, and W_t is

the wage of unskilled labor. The associated first-order conditions are:

$$\gamma \left(\frac{L_t}{M_t} \right)^{1-\gamma} = \frac{P_t^M}{P_{t,m}}, \qquad (1-\gamma) \left(\frac{M_t}{L_t} \right)^{\gamma} = \frac{W_t}{P_{t,m}}. \tag{3.5}$$

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:

$$M_{t,j} = f(Y_{t,m}) = Y_{t,m}. \tag{3.6}$$

The intermediate input producer operates in a monopolistically competitive market, facing a demand function characterized by (3.2). 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 $\Phi_{t,j}$:

$$\Phi_{t,j} = (P_{t,j}^M - P_{t,m}) \left(\frac{P_{t,j}^M}{P_t^M} \right)^{\frac{\vartheta}{1-\vartheta}} M_t. \tag{3.7}$$

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.

3.2 The Innovation Block

The innovation block in our model builds upon Romer (1990). There is a continuum of a unity measure of innovators. Each innovator *creates* new varieties of intermediate inputs and profits by selling the Intellectual Property Rights.

Figure 18 provides an overview of the innovation block, which encompasses three phases within each period. In Phase I, innovators receive new raw ideas and decide whether to initiate a new R&D project, utilizing the stock of raw ideas they have already accumulated. Following this, Phase II involves the random adoption of ideas embedded in existing R&D projects. Henceforth, ideas embedded in R&D projects are referred to as R&D ideas. This phase is characterized by the hiring of skilled labor for idea adop-

tion, with R&D expenses comprising the labor costs of this skilled labor. Finally, in Phase III, the successfully adopted R&D ideas materialize into new varieties of intermediate input. The innovator then monetizes these developments by selling the Intellectual Property Rights of these new varieties to the intermediate input producer described above in Section 3.1.

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 the paper is to consider innovators' extensive margin decision in Phase I: whether or not to set up a new R&D project that transforms raw ideas into R&D ideas. Next, we will describe the details of each phase.

Phase I: R&D Decisions Along the Extensive Margin We start with the central focus of the paper: R&D decisions along the extensive margin. In Phase I, innovators decide whether to initiate R&D projects, embedding their stock of raw ideas in preparation for adoption.

There is a continuum of heterogeneous innovators, each receiving an exogenous amount $\eta_{i,t}$ of raw ideas inflow every period, where $\eta_{i,t}$ is drawn from an independent and identically distributed distribution. The distribution is homogeneous across innovators. This heterogeneous inflow of raw ideas gives rise to a heterogeneous stock of raw ideas $s_{i,t}$. Throughout the paper, $s_{i,t}$ can be understood as either the amount of raw ideas or the quality of an idea; the two interpretations are isomorphic.

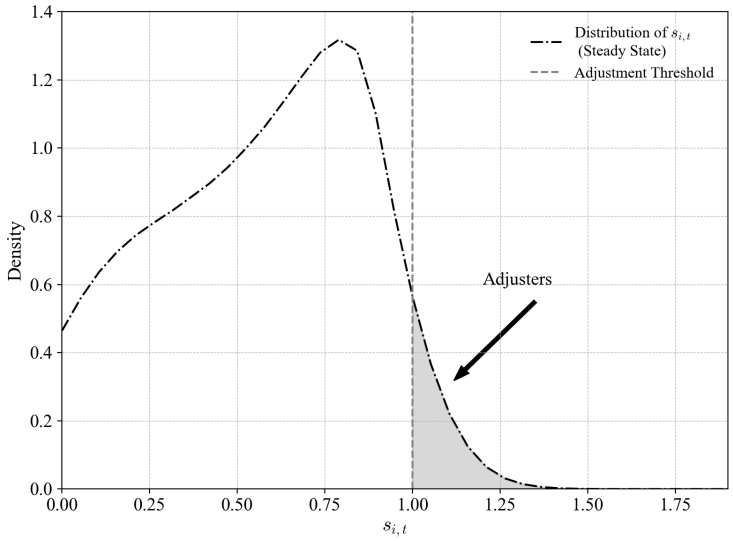
The innovator decides whether or not to allocate its unused raw ideas to a new R&D project. Importantly, there is a fixed cost for setting up a new R&D project, which amounts to cA_t. The fixed setup cost is proportional to the number of varieties A_t, reflecting that the cost of setting up an R&D project increases with the complexity of technology. With no convex adjustment cost, an innovator assigns all of their raw ideas to a single project once the decision to set up an R&D project is made. The unconverted raw ideas depreciate at the rate $1 - \phi$.

Denoting $S_t(s_{i,t})$ as the value of innovator i with idea stock $s_{i,t}$ at period t, the innovator's problem can be written as:

$$S_t(s_{i,t}) = \max\left\{S_t^a(s_{i,t}), S_t^n(s_{i,t})\right\}, \tag{3.8}$$

which reflects that innovators can decide whether to set up a new R&D project or not. Let $\tilde{J}_t$ denote the unit value of ideas, which will be further discussed in Phase II. The value of

Figure 4: R&D Decisions Along the Extensive Margin



an innovator i who decides to set up a new project is:

$$S_t^a(s_{i,t}) = -cA_t + s_{i,t}\tilde{J}_t + E_t\Lambda_{t,t+1}S_{t+1}(\eta_{i,t+1}). \tag{3.9}$$

It depends on $-cA_t$, which is the fixed setup cost, $s_{i,t}\tilde{J}_t$, the total value of R&D ideas, and $E_t\Lambda_{t,t+1}S_{i,t+1}(\eta_{i,t+1})$, the value of innovator i in the next stage starting with no idea stock.

The continuation value of an innovator i who decides to wait and not set up a new project in the current period is:

$$S_t^n(s_{i,t}) = E_t\Lambda_{t,t+1}S_{t+1}(\phi s_{i,t} + \eta_{i,t+1}). \tag{3.10}$$

Solving innovators' problem, as described in Equations (3.8) - (3.10), gives rise to the following proposition.

Proposition 1. *If $c > 0$, there exists a threshold $\bar{s}_t > 0$ such that innovators choose to set up a new R&D project if and only if $s_{i,t} > \bar{s}_t$.*

Proof. See Appendix D.3. □

Figure 4 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 amount of adjusters—innovators who decide to set up new projects—in the economy.

It is evident that the overall level of innovation in the economy depends on the distribution of raw ideas and $\bar{s}_{i,t}$. Business cycle shocks impact both factors, consequently influencing aggregate innovation.

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

$$s_{i,t+1} = \phi s_{i,t} I\{s_{i,t} < \bar{s}_{i,t}\} + \eta_{i,t+1}, \tag{3.11}$$

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

Phase II: The Adoption of Ideas Embedded in R&D Projects Once a new R&D project is decided, it becomes a pending project awaiting construction. In each period, a proportion $\tilde{\lambda} < 1$ of pending projects are realized and commence construction. This mechanism is designed to capture the reality that R&D projects do not materialize instantaneously; it also reflects the 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 be the unit value of R&D ideas (i.e., ideas in on-going R&D projects), $\Lambda_{t,t+1}$ represent the stochastic discount factor of the household, $1 - \phi_A$ denotes the obsolescence rate of technology, the unit value of an idea within a pending R&D project is expressed as:

$$\tilde{J}_t = \tilde{\lambda} J_t + \phi_A (1 - \tilde{\lambda}) E_t \Lambda_{t,t+1} J_{t+1}. \tag{3.12}$$

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

To ensure an adoption rate of λ, the innovator must hire a fixed amount $\bar{n}$ unit of skilled labor. The existing technologies have a spillover effect, where each unit of adopted R&D ideas will transform into A_t units of new varieties. R&D ideas that have not been 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:

$$J_t = -W_t^{RD} \bar{n} + \lambda E_t \Lambda_{t,t+1} A_{t+1} V_{t+1} + (1 - \lambda) \phi_A E_t \Lambda_{t,t+1} J_{t+1}, \tag{3.13}$$

where V_t denotes the unit value of a new variety that will be determined in Phase III, W_t^{RD} denotes the wage of skilled labor, $\bar{n}$ is the fixed amount of skilled labor.

The first term of (3.13) represents the cost of skilled labor. The second term is the value

of the R&D ideas that are successfully converted into new varieties of intermediate input. The third term is the expected value of R&D ideas that remains to be adopted.

Phase III: Monetizing the Intellectual Property Rights Once adopted, an R&D idea on a 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 of production for a single period. The value of the patent is defined as

$$V_t = \Phi_t, \tag{3.14}$$

where Φ_t is defined in Equation (3.7).

To recap, the value of the patent V_t, modeled in Phase III, determines the unit value of an idea in an R&D project, which ultimately influences innovators' extensive margin decisions on whether to set up a new project in Phase I. Next, we derive the evolution of the aggregate number of varieties A_t.

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 3.12, 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 $\tilde{x}^{new}_{i,t}$, is:

$$\tilde{x}^{new}_{i,t} = s_{i,t} I\{s_{i,t} \geq \bar{s}_{i,t}\}. \tag{3.15}$$

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:

$$\tilde{X}_{t+1} = \phi_A(1-\tilde{\lambda})\tilde{X}_t + \int_i \tilde{x}^{new}_{i,t+1}\, di, \tag{3.16}$$

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

$$X_{t+1} = \phi_A(1-\lambda)X_t + \tilde{\lambda}\tilde{X}_{t+1}. \tag{3.17}$$

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

$$A_{t+1} = \phi_A A_t + \lambda X_t A_{t+1}. \tag{3.18}$$

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 $\tilde{x}_{i,t}^{new}$ through Equation (3.15). This, in turn, ultimately affects the aggregate productivity A_t through Equations (3.16) to (3.18). The next section discusses the business cycle block.

3.3 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

$$\max_{\{C_{t+i},L_{t+i},N_{t+i},Q_{t+i}\}} E_t \sum_{i=0}^{\infty} \beta^i \varrho_{t+i} \left[\frac{C_{t+i}^{1-\sigma}}{1-\sigma} - \frac{L_{t+i}^{1+\varphi}}{1+\varphi} - \left(\frac{N_{t+i}^{1+\eta}}{1+\eta} \right) \right] \tag{3.19}$$

subject to

$$P_t C_t = W_t L_t + W_t^{RD} N_t + Q_t - R_t^{-1} Q_{t+1} + \Xi_t \tag{3.20}$$

where ϱ_t is a demand shock given by the following exogenous process: $\log \varrho_{t+1} = \rho_\varrho \log \varrho_t +$

$\sigma_\varrho e^\varrho_{t+1}$, with $e^\varrho_{t+1} \sim N(0,1)$. Solving the household's problem leads to the consumer Euler equation:

$$1 = E_t \left[\Lambda_{t,t+1} \Pi^{-1}_{t+1} R_t \right] \tag{3.21}$$

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

$$C^\sigma_t L^\varphi_t = \frac{W_t}{P_t}, \qquad C^\sigma_t N^\eta_t = \frac{W^{RD}_t}{P_t}. \tag{3.22}$$

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:

$$Y_t = \left(\int y^\mu_{t,h} \right)^{\frac{1}{\mu}} \tag{3.23}$$

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 $P_t = \left(\int p^{\frac{1}{1-\mu}}_{t,h} \right)^{1-\mu}$. The final good producer's cost minimization leads to the following demand function:

$$y_{t,h} = \left(\frac{p_{t,h}}{P_t} \right)^{\frac{\mu}{1-\mu}} Y_t. \tag{3.24}$$

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.

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:

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

subject to the demand constraint (3.24). $\Lambda_{t,t+j} \equiv \prod_{k=0}^{j} \Lambda_{t+k,t+k+1}$. 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):

$$\pi_t = \beta E_t \pi_{t+1} + \frac{\mu}{\kappa(1-\mu)} \left(\frac{1}{\mu} - \frac{P_{t,m}}{P_t} \right). \tag{3.25}$$

Monetary policy Following Cúrdia et al. (2015) and Zhang (2022), we assume that the central bank determines the interest rate by adopting a version of the Taylor rule that takes into account the efficient real interest rate, the one that would prevail if markets are free from distortion (Wicksell 1898):

$$R_t = R_t^e \Pi_t^{\varphi_\pi}, \tag{3.26}$$

In our framework, this efficient real interest rate emerges under the assumption of a competitive retail market[4]. The efficient real interest rate, R_t^e, is then implied by the equation of optimal intertemporal substitution:

$$1 = \beta E_t \Big[\Big(\frac{\varrho_{t+1}}{\varrho_t} \Big)^{\omega - 1} \frac{A_t}{A_{t+1}} R_t^e \Big] \tag{3.27}$$

where the coefficient ω is endogenously determined by the model parameters.

Discussions of Model Assumptions Appendix D.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 dis-

[4]In our model, the monopolistic power of the intermediate input producer is assumed to be non-distortionary, as it incentivizes innovation.

tribution 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.[5] 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[6], 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.

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

3.4 Market Clearing

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

$$Y_t = C_t, \qquad Y_{t,m} = \int_0^1 y_{t,h} dh + \int_0^{A_t} M_{t,j} dj. \tag{3.28}$$

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 $Q_t = 0$. Appendix D.2 defines the model's equilibrium, summarizes the equations characterizing it, and describes the solution method.

4 Model Calibration and Estimation

We calibrate the model to the U.S. economy at an annual frequency. Section 4.1 presents the calibration strategy, while Section 4.2 discusses the calibrated parameters and the model fit. In Section 4.3, the estimation of shocks is presented, given the calibrated model.

[5] 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.

[6] As an alternative approach, employing a Real Business Cycle (RBC) model that incorporates an exogenously determined stochastic discounting factor can achieve a similar objective.

4.1 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.

Fixed Parameters Table 2 summarizes the set of fixed parameters. The discount factor β is set to 0.96 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. We normalize the parameter ω to 1 as it merely scales the identified volatility of the preference shock without affecting the model's quantitative results. Finally, 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).

Calibration Strategy Next, we discuss the strategy for the calibration of the key parameters: ν, 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.

We assume the inflow of ideas to be $\eta_{i,t} = \nu\overline{\eta}_{i,t}$, where $\overline{\eta}_{i,t}$ is drawn from a standard normal distribution truncated below zero to reflect that it is more likely to draw a small amount (lower quality) of ideas than a large amount (higher quality). As the inflow of ideas is the source of economic growth, we include the average growth of real GDP to inform the calibration of ν.

The lumpiness of R&D 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 max-mean ratio of the R&D expenses 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 $\bar{n}$.

4.2 Calibration Results

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 10%, which is higher than the rate for technologies already in use. The parameter $\tilde{\lambda} = 0.43$ indicates that, on average, an R&D project takes about 1.5 years from planning to completion. Lastly, $c = 0.42$ implies that the fixed adjustment cost of R&D is approximately 42% 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.

Untargeted Moments: Rationalizing Facts 2 In our internal calibration, we focused on matching micro-level moments that reflect the lumpy nature of R&D, such as the aver-

Table 1: Calibration: Fitted Parameters

Parameter	Description	Value	Moments	Data	Model
			$E[\Delta \log(\text{GDP})]$	0.028	0.028
ν	Intensity of idea inflow	0.12	Time Duration between spikes	8.33	8.93
c	Fixed R&D cost	0.42	Max-mean ratio R&D expense	2.73	2.71
ϕ	Idea obsolescence	0.10	AR(1) spike rate	0.48	0.48
$\tilde{\lambda}$	Pending R&D projects	0.43	AR(1) $\Delta \log(\text{GDP})$	0.36	0.36
$\bar{n}$	Skilled labor for adoption	1.51	$\sigma(\Delta \log(\text{GDP}))$	0.017	0.017
σ_ϱ	SD of preference shocks	0.018	Corr. Spike and $\Delta \log(\text{GDP})$	0.42	0.43
			Corr. R&D and $\log(\text{GDP})$	0.62	0.67

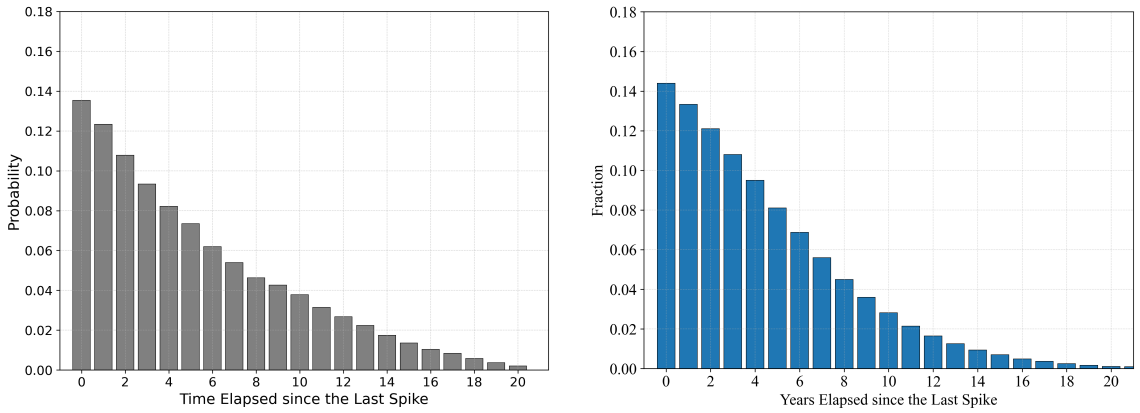
age duration between spikes and the mean-max ratio of R&D expenses. For untargeted moments, we concentrated on examining whether the model reproduces the aggregate characteristics of lumpy R&D. The last two rows of Table 1 report the moments in the data and the model's fit of the cyclicality of R&D and R&D spikes that are not targeted in the calibration. Our model closely matches the cyclicality of aggregate R&D. Due to the pro-cyclical nature of R&D project setup decisions and the unit labor cost of R&D, our model also captures the pro-cyclicality of the R&D spike rate.

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 5, 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 2 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 $s_{i,t}$.

Proposition 2. *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*

Figure 5: Distribution (Untargeted Moments): Time Elapsed since the Last Spike



(a) Data (Untargeted) (b) Model

Note: Figure 5 compares the firm-level distribution of *time elapsed since the last spike* between actual data and model simulations. For each firm, *time elapsed since the last spike* is calculated as the time interval between the current period and the occurrence of the most recent extensive margin adjustment.

relationship holds:

$$\mathcal{G}_t(s) = \sum_{n \in \mathcal{N}} P(s|TE_{i,t} = n)\mathcal{T}_t(n). \tag{4.1}$$

Figure 5 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.

4.3 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 5. 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^o.[7]

Figure 19a shows the fitting results, where the model-derived GDP growth rate closely matches the actual rate. Additionally, Figure 19a includes 80% and 95% probability bands

[7]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.

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. Notably, the declines in GDP growth rate during 2008 and 2009 are contained within the 95% and close to the 80% probability band, suggesting that the 2008 Great Recession was not due to a surprisingly large negative shock but rather was well-embedded in our model's mechanism.

As an external validation, Figure 19b 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. 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.

Notably, 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 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.

5 Model Results

Before turning to the results, Section 5.1 presents a novel mechanism that drives the key results. Section 5.3 presents the main result of the paper, the Great Twins hypothesis. Section 5.3 introduces an additional prediction of the model: the recession trap.

5.1 Novel Mechanism: Vulnerable Moderation

5.1.1 The Asymmetric Effects of Shocks on the Distribution

A lower volatility at the business cycle frequency, i.e., moderation, renders the economy more vulnerable. The key to this result lies in a *novel mechanism*: the asymmetric effects of shocks on the distribution of raw ideas.

To see this, first consider a positive demand shock. The shock generates an economic boom, increasing the profitability of developing new varieties. This motivates innovators to more eagerly initiate new R&D projects, resulting in the leftward shift of the adjustment threshold, as indicated in Figure 20a based on the calibrated economy departing from the steady state. Consequently, this positive shock prompts more adjusters to invest their accumulated ideas into new R&D projects, effectively compressing the distribution of $s_{i,t}$ in the subsequent period, as demonstrated by the red dashed line in Figure 20b.

Conversely, a negative shock renders the development of new varieties less profitable. Under such circumstances, innovators tend to prefer waiting for the economy to recover from the downturn, leading to a rightward shift of the adjustment threshold (Figure 21a). This shift, as a consequence, results in fewer innovators engaging in initiating new R&D projects during the economic downturn. Consequently, more innovators further accumulate their stock of innovative ideas. This accumulation process ultimately stretches the distribution of $s_{i,t}$, as shown in the green dashed line in Figure 21b.

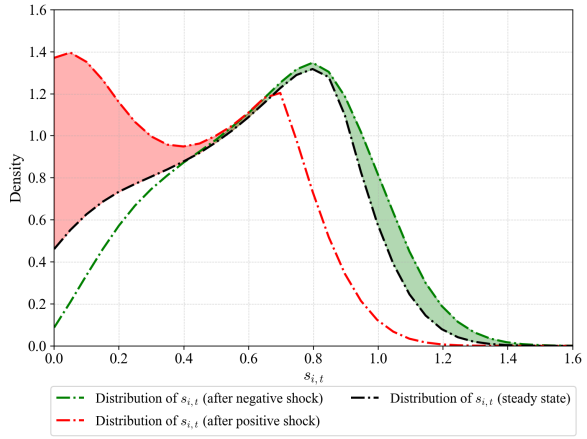
The stretched distribution due to a negative shock and the compressed distribution due to a positive shock are shown in Figure 6a. The following proposition shows that, under very general conditions, the compression from positive shocks (red shaded area) exceeds the stretching from negative shocks (green shaded area), given the presence of a fixed R&D cost.

Proposition 3. *For a given process of idea inflow $\eta_{i,t}$, there exists a threshold $\bar{c} > 0$, such that a positive shock to the adjustment threshold has a greater absolute effect on the distribution than a negative shock of the same size if and only if the adjustment cost $c > \bar{c}$.*

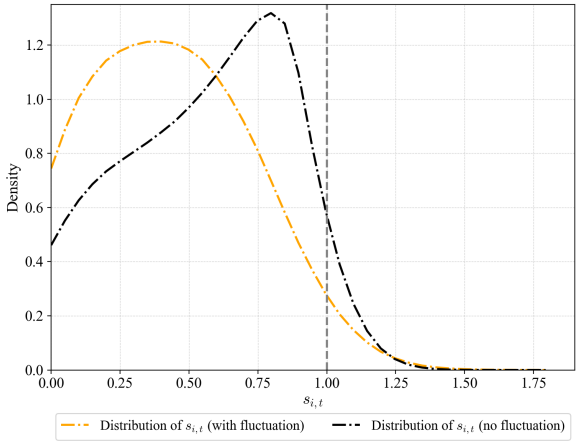
Proof. See Online Appendix D.4. □

The intuition behind this is as follows. For the asymmetric effects of positive and negative shocks on the idea distribution described above to be true, the adjustment threshold, $\bar{s}$, needs to be on the declining region of the steady state distribution of ideas. Intuitively,

Figure 6: The Asymmetry and Vulnerable Moderation



(a) Asymmetric Effects on the Distribution of $s_{i,t}$



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

this requires $\bar{s}$ to be sufficiently large, such that there is a diminishing probability of further accumulating ideas beyond the adjustment threshold. And because $\bar{s}$ is increasing in the fixed adjustment cost c, the fixed adjustment cost is essential for this result. 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.

In addition to c, another matrix that is critical for the key mechanism is the idea inflow distribution. Intuitively, if innovators received infinitely large new ideas each period, the adjustment threshold $\bar{s}$ would not fall in the declining region of the steady-state distribution of ideas. Proposition 3 considers a general process of idea inflow. The subsequent Corollary 1 demonstrates that, given a realistic density function of the idea inflow, a positive shock to the adjustment threshold has a greater absolute effect on the distribution than a negative shock, *for any* positive c.

Corollary 1. *The threshold $\bar{c}$ described in Proposition 3 equals zero if the density function of the idea inflow, $\psi(\eta)$, decreases with η.*

Proof. See Online Appendix D.4. □

Corollary 1 indicates that asymmetry is inherently present 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.

To further illustrate the critical role of lumpiness in deriving the model mechanism, in Online Appendix D.5, we define the degree of asymmetry as the difference between the red-shaded and green-shaded areas depicted in Figure 6a. We demonstrate analytically that, for two specific idea inflow processes with analytical solutions for the steady-state distribution of $s_{i,t}$, the degree of asymmetry increases with the magnitude of fixed adjustment costs (degree of lumpiness). Furthermore, our numerical exercise confirms that this pattern is consistent across various idea inflow processes and model parameters.

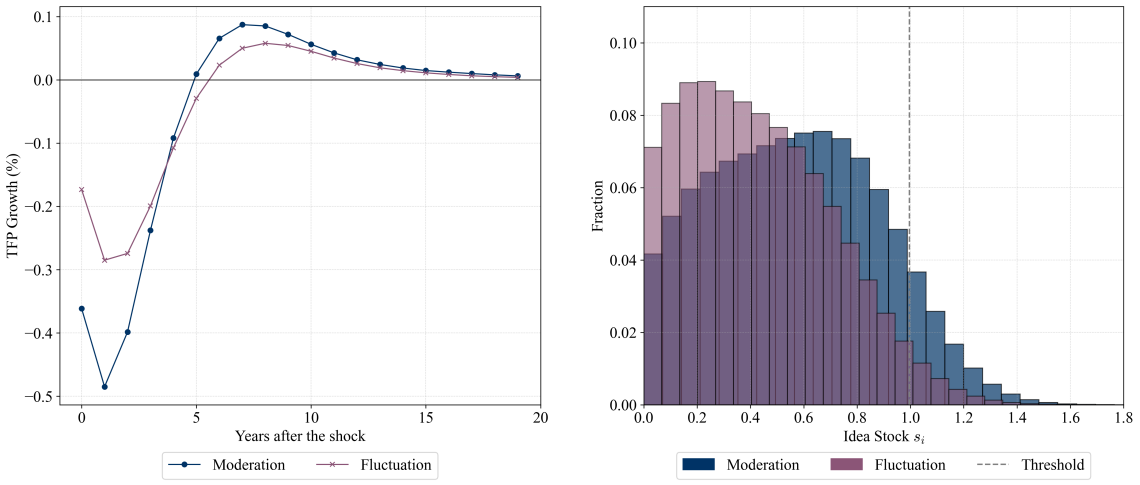
Discussion: Relationship with Lumpy Capital Investment Models Online Appendix D.6 discusses whether the same mechanism holds in a lumpy capital investment model. Briefly, since capital adjustments can occur in either direction, aggregate shocks do not lead to a stretching or compressing of the distribution emphasized above. By calibrating the parsimonious lumpy investment model to Compustat data, we show that the influence of aggregate fluctuations on the distribution of the capital gap is negligible.

5.1.2 The Distribution depends on History of Shocks

Now consider an economy with both positive and negative shocks. Due to the asymmetric effects of these shocks highlighted above, the net effects lead to a compression of the distribution of $s_{i,t}$, as illustrated in Figure 6b. In this figure, the orange dashed line represents the average distribution of $s_{i,t}$ given economic fluctuations, while the black dashed line depicts the steady state distribution of $s_{i,t}$, representing an extreme scenario where there is an absence of any fluctuations. Notably, with business cycle fluctuations, there is a considerable reduction in mass around the adjustment threshold.

Figure 7 illustrates the effect of moderation in our calibrated model. The moderation (fluctuation) states are states that have a *sample variance* in the last 10 periods smaller (larger) than the average. As suggested by the exposition of the mechanism, the distribution of $s_{i,t}$ under moderation states has a smaller mass around the extensive margin, comparing to the distribution of $s_{i,t}$ under fluctuation states, as shown in Figure 7b. Consequently, the impulse response of the growth of A_t to a one-sigma negative demand shock yields a larger effect under a moderation state than under a fluctuation state, as shown in Figure 7a.

Figure 7: The Impulse Response of Shocks in the Calibrated Model



(a) State-Dependent Effects of Shocks (b) Mechanism: Distribution & Threshold

Note: Figure 7a and 7b are based on simulations. We simulate 100,000 panels, and each panel contains 120 periods and 5,000 innovators. The moderation (fluctuation) subsample is defined as panels that have the lowest (highest) 5% standard deviation of output growth from $t = 91$ to $t = 100$. Figure 7b depicts the average distribution of $s_{i,t}$ in period $t = 100$. We re-simulated the subsamples of moderation and fluctuation with an additional one-standard-deviation negative demand shock imposed at period $t = 101$. Figure 7a plots the average impulse response of the moderation and fluctuation subsamples.

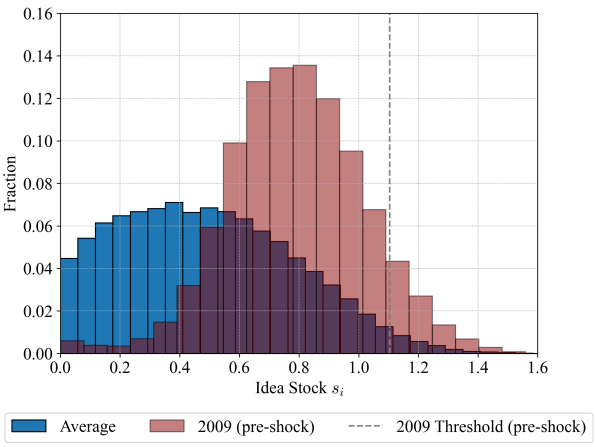
5.2 Model Results: the Great Twins

The results presented in Figure 7 correspond to what occurred during the Great Recession, according to the estimated model. We take a few steps to verify this. First, we retrieved the distribution of $s_{i,t}$ prior to the 2009 shock, which was the biggest negative shock in the sample given the annual frequency of the model (see Figure 19a).

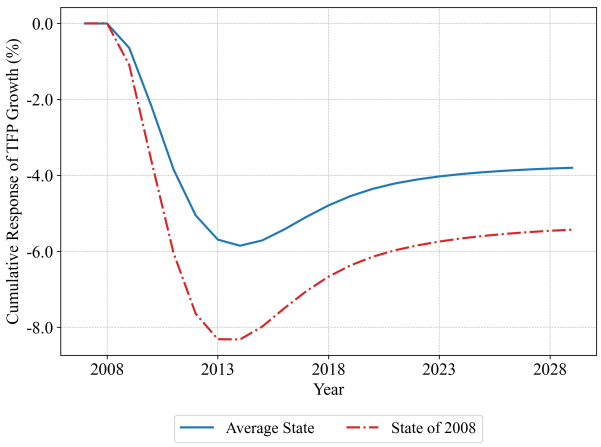
Figure 8a plots the pre-shock distribution of idea stock in 2009 and an average distribution of idea stock. We observe that the 2009 pre-shock distribution of idea stock, compared to the average, has a substantial mass concentrated around the adjustment threshold, implying a considerable risk of loss in R&D activities.

Figure 8b illustrates the decline in A_t resulting from the 2008 and 2009 recessionary shocks. In comparison to the direct impact of the demand shock, the effects of the Great Recession shocks on A_t unfolded gradually, reaching their peak around 2014 with a slight rebound in the subsequent years. This rebound can be attributed primarily to innovative ideas that were not invested in during the recession but accumulated over time, eventually materializing into new technologies.

Figure 8: The Distribution of Idea Stock before GFC and the CIR of of TFP to the GFC shocks



(a) Distribution before the Great Recession



(b) CIR of TFP to the Great Recession Shocks

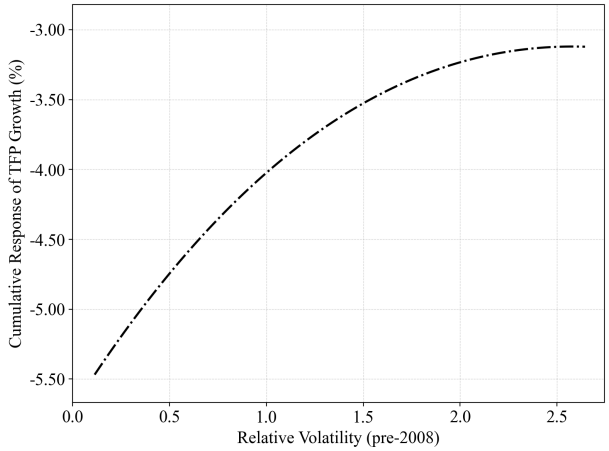
Note: Figure 8a compares the average distribution of $s_{i,t}$ with its 2009 distribution, based on the model fitted to actual GDP growth. Figure 8b shows the cumulative TFP response to the 2008 and 2009 shocks, contrasting the average state with the state prior to the 2008 shock. The average state is defined by having the same ϱ_t as the compared state, though the distributions of $s_{i,t}$ may differ.

The results of the Cumulative Impulse Response (CIR) of A_t indicate that the U.S. economy, having experienced the Great Recession, now exhibits a GDP nearly 5.8% lower than a counterfactual U.S. economy without the Great Recession shocks.[8] More importantly, 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, averaging around 3.9%.

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 assert that this pre-shock distribution, implying a notable left-tail risk, was caused by the moderated economic fluctuations during the 2003-2007 period. To investigate this, we simulated a counterfactual sample of the U.S. economy for 2003-2007. We calculated the ratio of estimated volatility to the average 5-period volatility in this sample, along with the corresponding loss in A_t. Figure 9 illustrates the relationship between 5-period relative volatility and the loss in A_t. We find that if the relative volatility increased from 0.5 to 1.5, the average loss in A_t caused by Great Recession shocks would be reduced by 1.5

[8]Note that this does not imply that the GDP is 5.8% lower than its pre-shocks level since the counterfactual U.S. economy without the Great Recession shocks would have experienced positive growth.

Figure 9: Great Twins: Effects of Great Recession Shocks with Varying Moderation



Note: Figure 9 show the cumulative impulse responses to the shocks of 2008 and 2009, varying by moderation levels, based on simulations. For the period 1979 to 2002, ϱ_t corresponds to the fitted values from Section 4.3. We simulate a counterfactual ϱ_t for 2003 to 2007. Relative volatility is defined as the ratio of the realized volatility of ϱ_t to its unconditional realized volatility for 2003-2007.

percentage points.

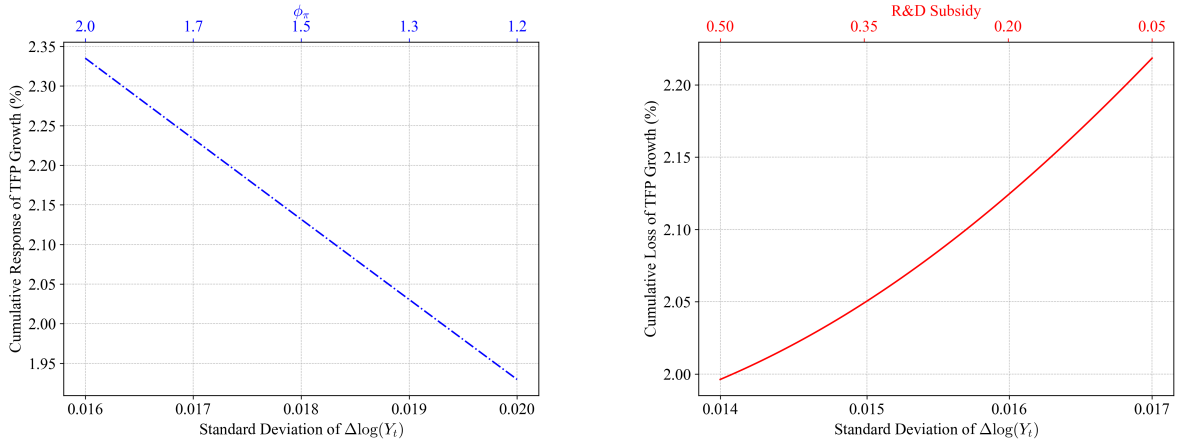
It is worth clarifying that reduced sample volatility does not imply reduced risk. Specifically, the changes along the x-axis in Figure 9 are entirely due to luck—more precisely, the Great Moderation resulted from drawing smaller shocks, while the risks—volatilities of shocks and unconditional volatility of the economy—remained constant.

Policy Implication I: The Tail-Risk Variance Trade-off In the previous simulations, the differences in sample volatilities were attributed to luck. This represents the 'Good Luck' view (Justiniano and Primiceri 2008). An alternative view is that the Great Moderation is a result of a 'Good Policy' (Galí and Gambetti 2009): the U.S. Federal 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 10a 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 seen

Figure 10: The Great Twins: Policy Implications



(a) The Tail-Risk Variance Trade-off

(b) Splitting the Great Twins

Note: Figures 10a and 10b illustrate the cumulative impulse responses to the shocks of 2008 and 2009, across different policy parameter settings. The top x-axis denotes the policy parameter values, and the bottom x-axis corresponds to the standard deviation of GDP growth. The y-axis measures the cumulative loss in TFP growth. Both figures are based on simulations. For each value of policy parameters, we simulate 5,000 panels, and each panel contains 140 periods and 5,000 innovators. The shocks happen in t=101 and t=102. The cumulative loss in TFP is calculated by averaging the cumulative impulse response of TFP in the last 10 periods.

in the correspondence between the upper and lower horizontal axes. The Y-axis represents the left-tail risk; specifically, the cumulative loss of TFP caused by Great Recession shocks. The downward-sloping blue dashed line indicates that reducing overall economic volatility through monetary policy adjustments can increase economic tail risk.

Policy Implication II: Splitting the Great Twins The Great Twins pose a novel dilemma for the Central Bank. Is there a complementary policy that can reduce overall economic fluctuations without increasing economic tail risk? The next exercise addresses this policy question.

We introduce a subsidy for the conduct of R&D. Specifically, the government offers a one-time lump-sum subsidy to innovators who start new R&D projects. This subsidy effectively reduces the fixed adjustment cost friction in the innovation block. We then evaluate the tail-risk and variance relationship in the model with different levels of R&D subsidy.

Figure 10b presents the results. Increasing R&D subsidies (upper horizontal axis) can reduce overall economic volatility (lower horizontal axis) while simultaneously mitigat-

ing the medium-term loss caused by Great Recession shocks (vertical axis).

5.3 Model Results: the Recession Trap

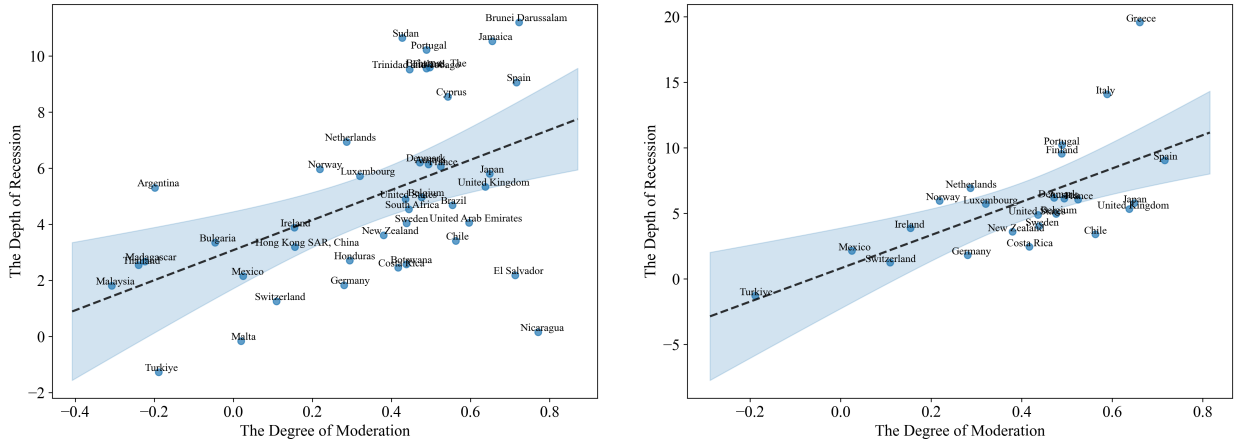
Our model also implies another feature related to economic policy: the effectiveness of nominal demand stimulus policy diminishes during recessions—a phenomenon we refer to as the Recession Trap. Figure 23a presents the impulse response of the growth rate of A_t to a 25 basis points monetary shock. To conduct this policy analysis, we have introduced exogeneous monetary policy shocks ε_t^m to the Taylor rule, with $\varepsilon_{t+1}^m = \rho_m \varepsilon_t^m + e_{t+1}^m$.

Figure 23a shows that the effects of monetary policy during recessions, periods where the effects of stimulus policy are most appreciated, are substantially dampened compared to the effects in a boom state. At their peak, the effects in a recession state are reduced to one-third of the effects in a boom state.

Figure 23b explains the reason behind this implication. The result is mainly due to the difference in the adjustment threshold. During a recession, the adjustment threshold is far to the right compared to its counterpart during a boom. This is because, in a recession, the profit from producing new varieties decreases, reducing the willingness to establish new R&D projects. Consequently, there is a relative scarcity of innovators moving to conduct new R&D projects during a recession state, thus diminishing the overall impact of nominal demand stimulus policies that shift the adjustment threshold to the left.

Policy Implication III: Escaping the Recession Trap Is there a policy that can enhance the effectiveness of monetary stimulus during a recession? Yes, if a monetary policy shock is introduced together with a one-time R&D subsidy. Figure 22 presents the result by plotting the impulse response function of a monetary shock in a recession state under alternative scenarios. The red dash-dotted line shows the net effects (impulse response function) of a monetary shock when combined with an R&D subsidy. Compared to scenarios without the subsidy (the blue solid line), this R&D subsidy enhances the effectiveness of monetary shocks during recessions. The primary reason is that the R&D subsidy reduces the idea stock needed to reach the extensive margin. Consequently, even in a recession, the mass of innovators near the extensive margin is higher than in scenarios without the subsidy, thereby strengthening the stimulative effect of the monetary shock on the economy.

Figure 11: The Great Twins in the Data



(a) Full Set

(b) OECD Countries

Notes: Figure 11a plots the degree of moderation against the depth of recessions for countries from the World Bank Database. We consider countries with at least 25 observations before 2008, and tiny economies with a Real GDP of less than 10 billion dollars in 2019 are excluded. For Figure 11a, we trim the data of the depth of recession by using the Median Absolute Deviation approach with parameter $c = 3$. The pre-recession growth trend is defined as the economy continuing to grow at the average growth rate of the past fifteen years. 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. Figure 11b reports the same relationship for the 38 OECD countries. The correlation in Figure 11a is 0.50 ($p = 0.00$), and the correlation for the OECD countries is 0.64 ($p = 0.00$).

6 Empirical Validation of the Model Predictions

Next, we return to the data to seek evidence that supports the model's predictions and its underlying mechanisms. Section 6.1 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 6.2 presents micro-level evidence that supports this mechanism.

6.1 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-

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.[9]

Figure 11a shows that the depth of recessions positively correlates with the degree of moderation in the pre-crisis periods, consistent with the Great Twins hypothesis. Figure 11b demonstrates the same result by restricting the sample to OECD economies.

Additional Results In Online Appendix C.4, we further show the relationship holds when financial factors are controlled for. Our model attributes the depth of the Great Recession to the drop in TFP. Figure 24 presents the relationship between the TFP-measured depth of the recession and the degree of moderation. Next, we ask whether the model-implied relation between moderation and recession is also evident in recessions beyond the 2008 Great Recession. To this end, we repeat the previous exercise using data that covers 18 advanced economies since World War II. Figure 25 shows a positive correlation between the depth of recessions and the degree of moderation for 18 advanced economies from 1950 to 2019. We conclude that the Great Twins is a robust phenomenon across countries and over time.

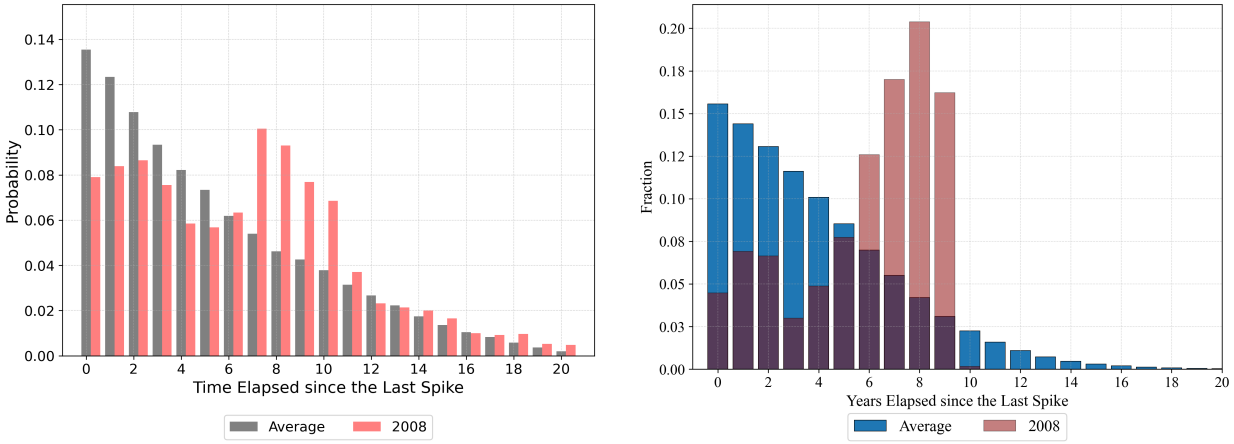
6.2 Micro-Evidence: Distribution in the Data

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 8a, is crucial for the observed depth of the Great Recession. In Section 4.1, 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 2, 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 12 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 12b, is due to a large number of innovators being 6-9 years past their last R&D spike, coinciding with the average 8.33-year gap between R&D spikes in the data. As shown in Figure 12a, this pattern

[9] See Appendix C.4 for details of the formal definition of the depth of recession.

Figure 12: Distribution of Time Elapsed since the Last Spike: Pre-2008 Recession



(a) Distribution: in the Data

(b) Distribution: in the Model

Note: Figure 12a 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. Figure 12b 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.

is also mirrored in the empirical data.

7 Conclusion

We document three novel facts. First, R&D investment decisions are lumpy: innovative firms often do not adjust R&D expenses, but when they do, they adjust by a substantial amount. Second, the fraction of firms making spike adjustments in R&D expenses—the spike rate—is pro-cyclical. Third, the pro-cyclicality of these spike adjustments explains the pro-cyclicality of aggregate R&D expenses as documented in the existing literature.

Motivated by these facts, we develop a medium-run business cycle model with heterogeneous innovators facing a fixed setup cost for new R&D projects that gives rise to R&D decisions along the extensive margin. We highlight the role of the distribution of undeveloped innovative ideas among innovators. Specifically, we illustrate that a new mechanism—the asymmetric effects of business cycle fluctuations on the distribution—gives rise to the Great Twins: The Great Moderation caused the observed severity of the subsequent Great Recession. Moreover, we discuss policy implications and a 'recession trap' in the model—nominal demand stimulus policies are muted during recessions.

Last but not least, we document the Great Twins in the data, and we uncover micro-level evidence supporting the underlying mechanism.

References

D. Acemoglu, P. Aghion, L. Bursztyn, and D. Hemous. The environment and directed technical change. *American economic review*, 102(1):131–166, 2012.

D. Acemoglu, U. Akcigit, H. Alp, N. Bloom, and W. Kerr. Innovation, reallocation, and growth. *American Economic Review*, 108(11):3450–91, 2018.

P. Aghion and P. Howitt. A model of growth through creative destruction. *Econometrica*, 1992.

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

D. Anzoategui, D. Comin, M. Gertler, and J. Martinez. Endogenous technology adoption and r&d as sources of business cycle persistence. *American Economic Journal: Macroeconomics*, 11(3):67–110, 2019.

A. Auclert. Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–2367, 2019.

R. Bachmann, R. J. Caballero, and E. M. R. A. Engel. Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model. *American Economic Journal: Macroeconomics*, 5(4):29–67, 2013. doi: 10.1257/mac.5.4.29.

F. C. Bagliano and C. Morana. It ain't over till it's over: A global perspective on the great moderation-great recession interconnection. *Applied Economics*, 49(49):4946–4969, 2017.

I. Baley and A. Blanco. Aggregate dynamics in lumpy economies. *Econometrica*, 89(3): 1235–1264, 2021.

L. Ball. Long-term damage from the great recession in oecd countries. *European Journal of Economics and Economic Policies: Intervention*, 11(2):149–160, 2014.

G. Barlevy. On the cyclicality of research and development. *American Economic Review*, 97 (4):1131–1164, 2007.

C. Bayer, R. Luetticke, L. Pham-Dao, and V. Tjaden. Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica*, 87 (1):255–290, 2019.

G. Benigno and L. Fornaro. Stagnation traps. *The Review of Economic Studies*, 85(3):1425–1470, 2018.

F. Bianchi, H. Kung, and G. Morales. Growth, slowdowns, and recoveries. *Journal of Monetary Economics*, 101:47–63, 2019.

N. Bloom. Uncertainty and the dynamics of r&d. *American Economic Review*, 97(2):250–255, 2007.

N. Bloom, C. I. Jones, J. Van Reenen, and M. Webb. Are ideas getting harder to find? *American Economic Review*, 110(4):1104–1144, 2020.

T. Broer, N.-J. Harbo Hansen, P. Krusell, and E. Öberg. The new keynesian transmission mechanism: A heterogeneous-agent perspective. *The Review of Economic Studies*, 87(1): 77–101, 2020.

J. R. Brown and B. C. Petersen. Cash holdings and r&d smoothing. *Journal of Corporate Finance*, 17(3):694–709, 2011.

R. J. Caballero and E. M. Engel. Explaining investment dynamics in us manufacturing: a generalized (s, s) approach. *Econometrica*, 67(4):783–826, 1999.

R. J. Caballero, E. M. Engel, J. C. Haltiwanger, M. Woodford, and R. E. Hall. Plant-level adjustment and aggregate investment dynamics. *Brookings papers on economic activity*, pages 1–54, 1995.

G. Cette, J. Fernald, and B. Mojon. The pre-great recession slowdown in productivity. *European Economic Review*, 88:3–20, 2016.

Z. Chen, Z. Liu, J. C. Suárez Serrato, and D. Y. Xu. Notching r&d investment with corporate income tax cuts in china. *American Economic Review*, 111(7):2065–2100, 2021.

J. Cloyne, C. Ferreira, M. Froemel, and P. Surico. Monetary policy, corporate finance, and investment. *Journal of the European Economic Association*, 21(6):2586–2634, 2023.

D. Comin and M. Gertler. Medium-term business cycles. *American Economic Review*, 96 (3):523–551, 2006.

D. A. Comin, M. Gertler, and A. M. Santacreu. Technology innovation and diffusion as sources of output and asset price fluctuations. Technical report, National Bureau of Economic Research, 2009.

R. W. Cooper and J. C. Haltiwanger. On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633, 2006.

V. Cúrdia, A. Ferrero, G. C. Ng, and A. Tambalotti. Has us monetary policy tracked the efficient interest rate? *Journal of Monetary Economics*, 70:72–83, 2015.

A. Fatas. Do business cycles cast long shadows? short-run persistence and economic growth. *Journal of economic growth*, 5:147–162, 2000.

J. Fernald. A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco, 2014.

J. Galí and L. Gambetti. On the sources of the great moderation. *American Economic Journal: Macroeconomics*, 1(1):26–57, 2009.

M. Gnewuch and D. Zhang. Monetary policy, firm heterogeneity, and the distribution of investment rates. *SSRN Working Paper*, 2023.

N. Gornemann, K. Kuester, and M. Nakajima. Doves for the rich, hawks for the poor? distributional consequences of monetary policy. 2016.

N. Gornemann, P. Guerrón-Quintana, and F. Saffie. Real exchange rates and endogenous productivity. 2021.

Z. Griliches. Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28(4):1661–1707, 1990.

P. A. Guerron-Quintana and R. Jinnai. Financial frictions, trends, and the great recession. *Quantitative Economics*, 10(2):735–773, 2019.

C.-T. Hsieh and P. J. Klenow. Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, 124(4):1403–1448, 2009.

Ò. Jordà, M. Schularick, and A. M. Taylor. Macrofinancial history and the new business cycle facts. *NBER macroeconomics annual*, 31(1):213–263, 2017.

J. Jungherr, M. Meier, T. Reinelt, I. Schott, et al. Corporate debt maturity matters for monetary policy. *Econometrica*, 1(4):345, 2022.

A. Justiniano and G. E. Primiceri. The time-varying volatility of macroeconomic fluctuations. *American Economic Review*, 98(3):604–641, 2008.

G. Kaplan, B. Moll, and G. L. Violante. Monetary policy according to hank. *American Economic Review*, 108(3):697–743, 2018.

A. Khan and J. Thomas. Idiosyncratic shocks and the role of nonconvexities. *Econometrica*, 76(2):395–436, 2008.

T. J. Klette and S. Kortum. Innovating firms and aggregate innovation. *Journal of political economy*, 112(5):986–1018, 2004.

Y. Koby and C. K. Wolf. Aggregation in Heterogeneous-Firm Models: A Sufficient Statistics Approach *. Technical report, 2020.

D. Krueger, K. Mitman, and F. Perri. Macroeconomics and heterogeneity, including inequality. *Handbook of Macroeconomics*, 2015.

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

P. Krusell, T. Mukoyama, and A. Şahin. Labour-market matching with precautionary savings and aggregate fluctuations. *The Review of Economic Studies*, 77(4):1477–1507, 2010.

H. Kung and L. Schmid. Innovation, growth, and asset prices. *The Journal of Finance*, 70 (3):1001–1037, 2015.

H. Lee. Striking while the iron is cold: Fragility after a surge of lumpy investments. *Mimeo, University of Tokyo*, 2023.

F. Lippi and A. Oskolkov. A structural model of asymmetric lumpy investment. 2023.

A. McKay and R. Reis. The role of automatic stabilizers in the us business cycle. *Econometrica*, 84(1):141–194, 2016.

G. Moscarini and F. Postel-Vinay. The cyclical job ladder. *Annual Review of Economics*, 10: 165–188, 2018.

P. Ottonello and T. Winberry. Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, 88(6):2473–2502, 2020.

A. Queralto. A model of slow recoveries from financial crises. *Journal of Monetary Economics*, 114:1–25, 2020.

M. Reiter, T. Sveen, and L. Weinke. Lumpy investment and the monetary transmission mechanism. *Journal of Monetary Economics*, 60(7):821–834, 2013.

P. M. Romer. Endogenous technological change. *Journal of political Economy*, 98(5, Part 2): S71–S102, 1990.

T. Schmitz. Endogenous growth, firm heterogeneity and the long-run impact of financial crises. *European Economic Review*, 132:103637, 2021.

P. Sedláček and V. Sterk. The growth potential of startups over the business cycle. *American Economic Review*, 107(10):3182–3210, 2017.

H. Tang and D. Zhang. Bubbly firm dynamics and aggregate fluctuations. *Journal of Monetary Economics*, 132:64–80, 2022.

H. Tang and D. Zhang. Bubbly credit cycles. *Available at SSRN 4285318*, 2023.

K. Wälde and U. Woitek. R&d expenditure in g7 countries and the implications for endogenous fluctuations and growth. *Economics Letters*, 82(1):91–97, 2004.

K. Wicksell. Interest and prices (english translation by rf kahn, 1936), 1898.

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

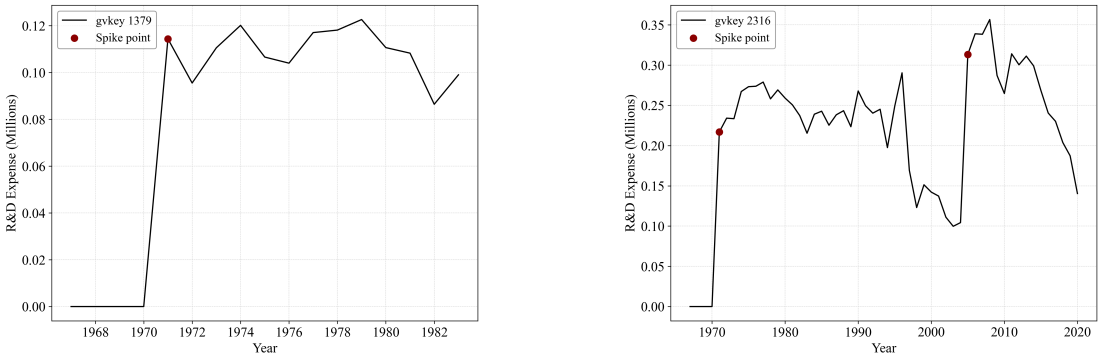
D. Zhang. Term structure, forecast revision, and the signaling channel of monetary policy. *Journal of the European Economic Association*, 20(4):1522–1553, 2022.

Online Appendix to "Lumpy R&D and the Great Twins: Great Moderation and Great Recession"

Lixing Wang, Donghai Zhang

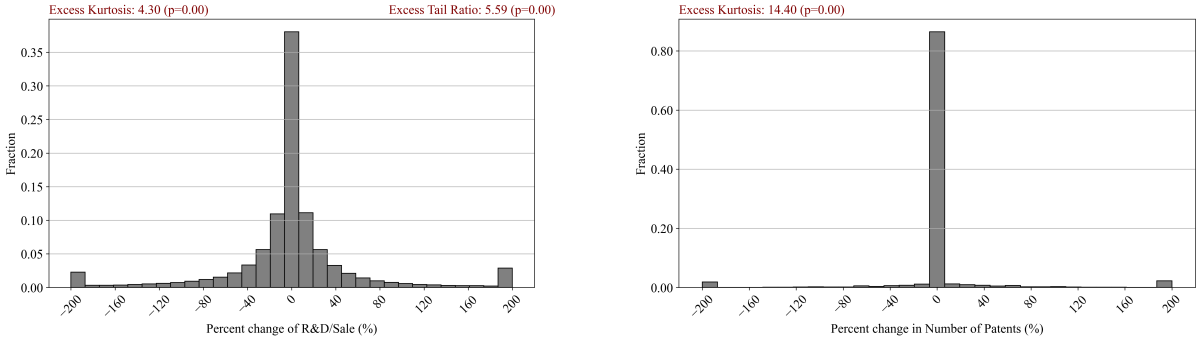
A Additional Figures

Figure 13: The Concept of Lumpy R&D Expenses



Note: This figure illustrates the evolution of R&D expenses for two actual firms in the data.

Figure 14: Fact 1 Robustness: R&D Metrics

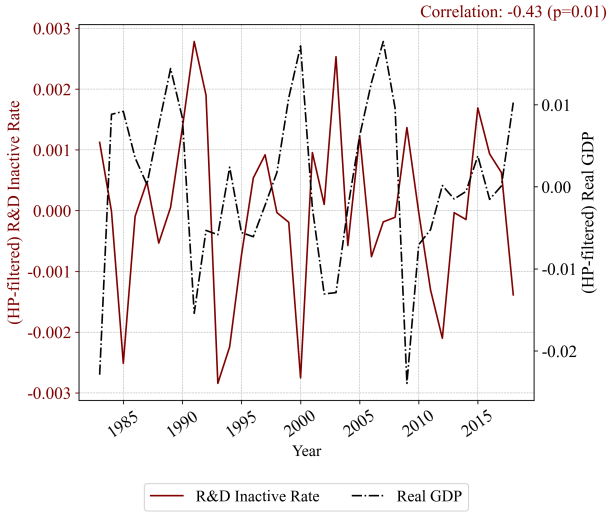


(a) Changes in the R&D-to-Sale Ratio

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

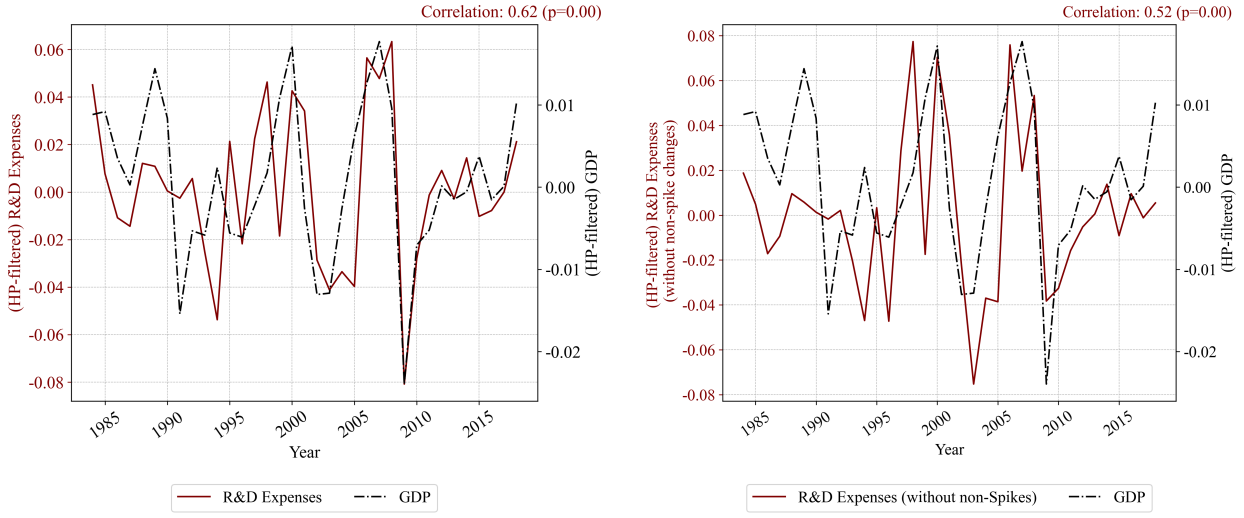
Note: The left figure depicts the histogram of changes in the R&D-to-sale ratio, excluding observations with non-positive sales. The right figure shows the histogram of changes in the number of new R&D patents annually, with data sourced from the USPTO. We select assignees with positive newly issued patents for at least five years.

Figure 15: Fact 2 Robustness: Cyclycality of Inaction



Note: The inactive rate is defined as the fraction of firms that do not change R&D expenses. Both series are filtered by HP filter, with $\lambda = 6.25$.

Figure 16: Fact 3 Robustness: R&D Spikes Drive the pro-cyclicality of Aggregate R&D



(a) Aggregate R&D (b) Aggregate R&D (without non-Spikes)

Note: The figures plot the time series of aggregate R&D spendings (with all and with only R&D spikes) and GDP. The red line is the R&D spendings and the blue line is the GDP. In figure 16b firms' R&D expenses are excluded from the aggregation if they are not identified as R&D spikes. We include positive and negative spikes to prevent artificially low correlations arising from excluding the pro-cyclical component. Both series are filtered by the HP filter, with $\lambda = 6.25$. The GDP data is from the FRED database.

Figure 17: Block Diagram of the Model System – A Comprehensive Overview

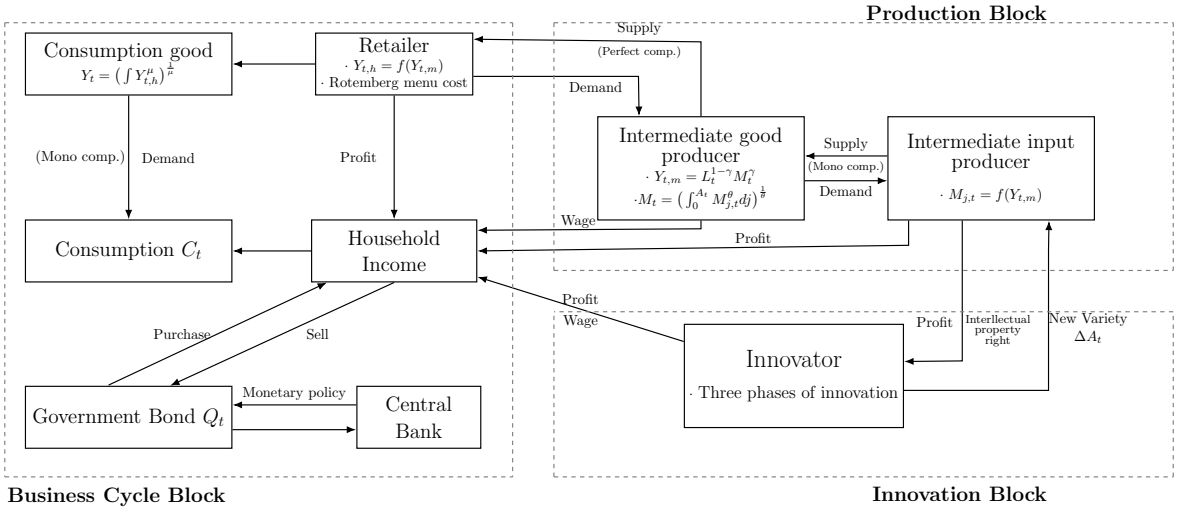


Figure 18: Block Diagram of the Model System – the Innovation Block

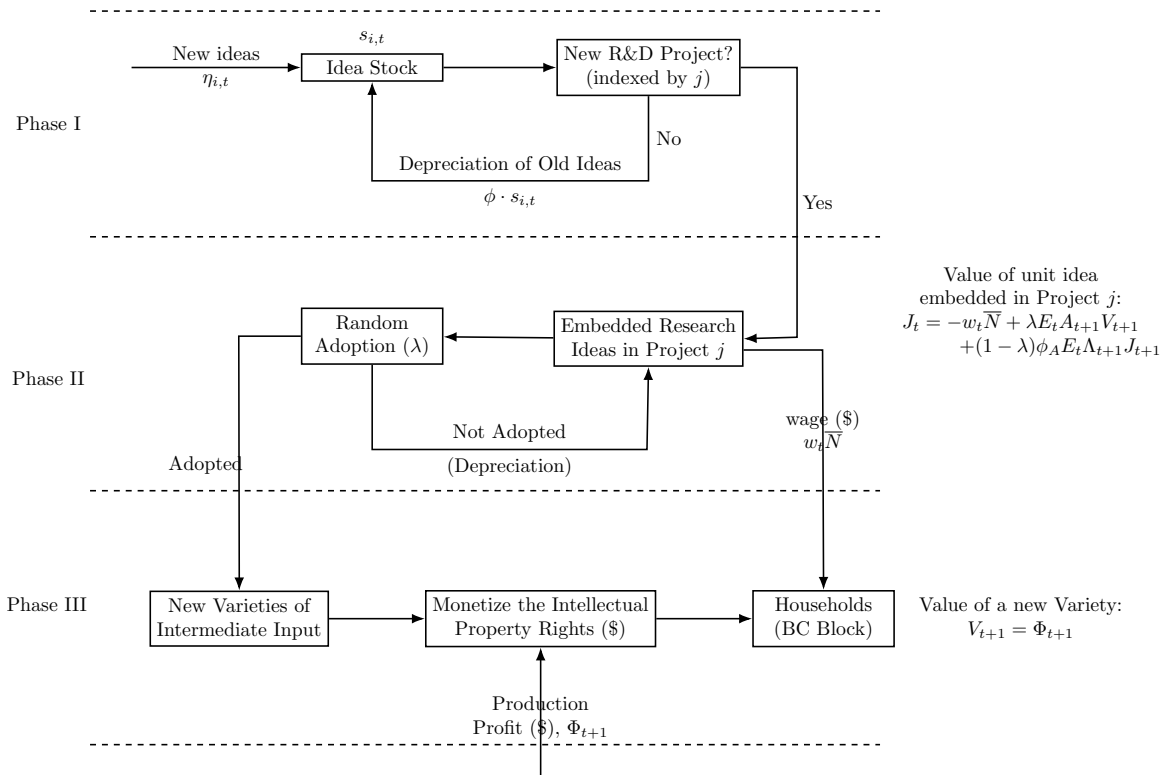
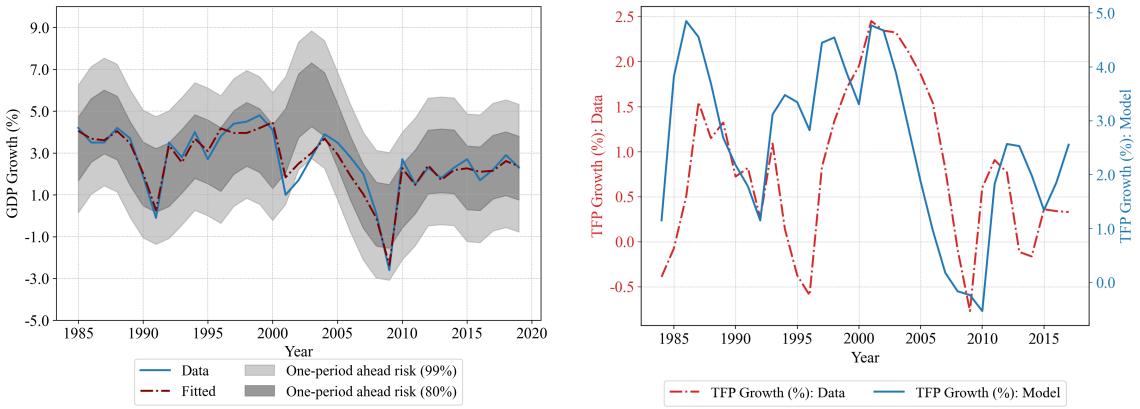


Figure 19: Fitting Model to the Data over Time

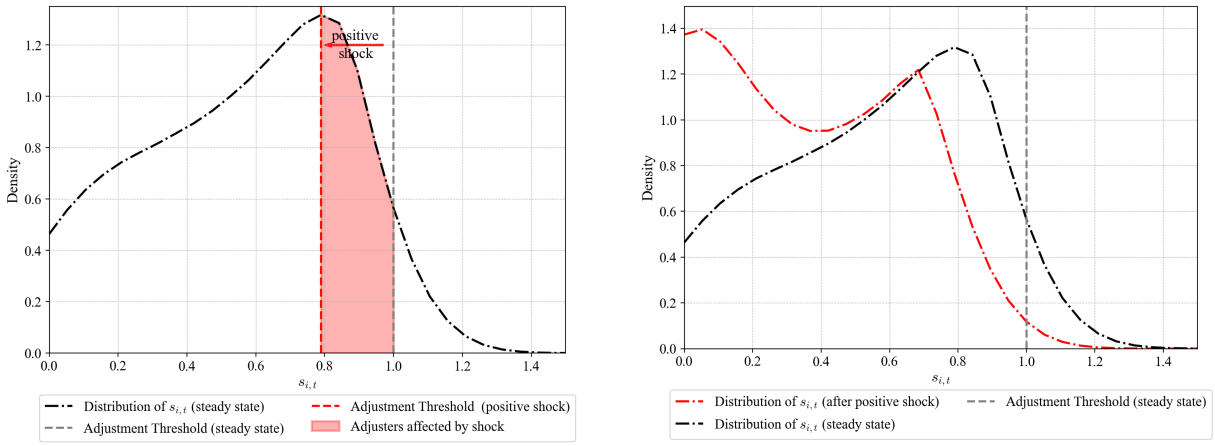


(a) Model Fit (in-sample): GDP

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

Note: Figure 19a reports the result of fitting model to the data. Figure 19b reports the model fitted TFP and actual TFP as an out-of-sample validation. The TFP data is from the Federal Reserve Bank of San Francisco, which report 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.

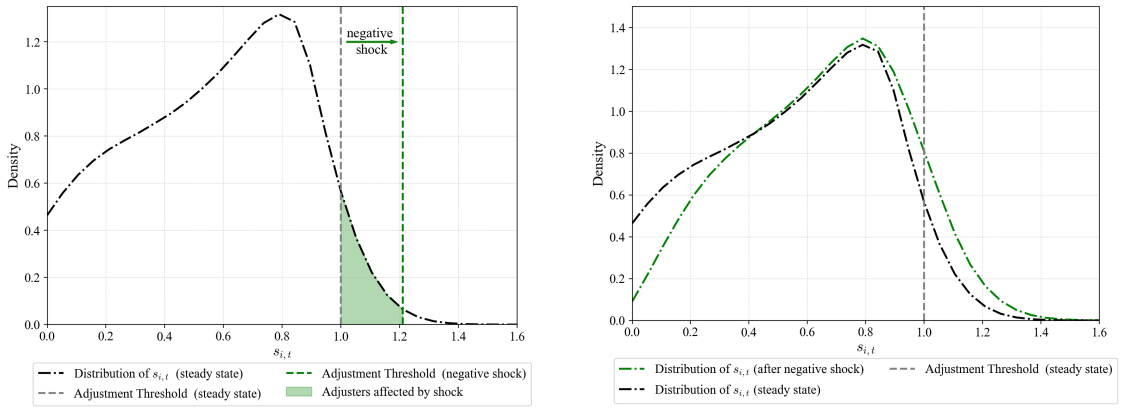
Figure 20: The Effect of a Positive Shock on the Distribution of $s_{i,t}$



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

(b) The distribution after a positive shock

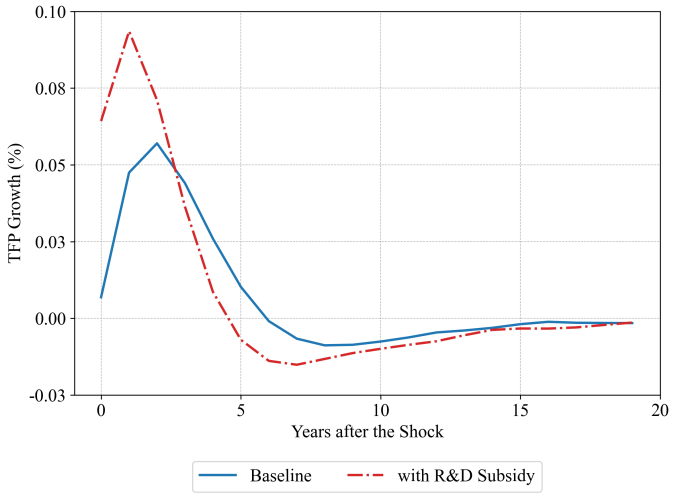
Figure 21: The Effect of a Negative Shock on the Distribution of $s_{i,t}$



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

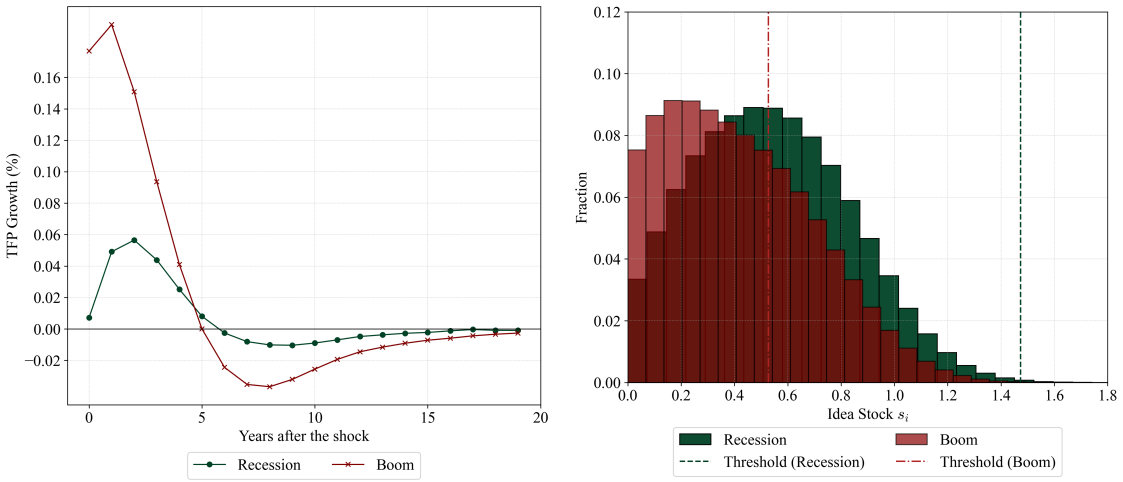
(b) The distribution after a negative shock

Figure 22: Escaping the Recession Trap



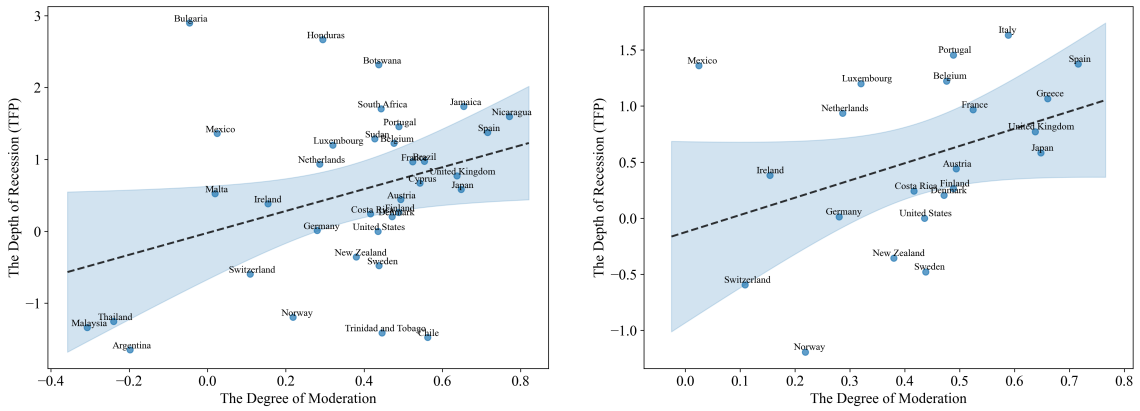
Note: Figure 22 is based on simulations. We simulate 100,000 panels, and each panel contains 120 periods and 5,000 innovators. Recession is defined as ϱ_t equals to $std(\varrho_t)$ at period $t = 100$. The blue solid curve represents the average impulse response to a monetary shock of 25 basis points in the recession state. The red dashed curve illustrates the response to a 25-basis-point monetary shock in the recession state, under the condition of an R&D subsidy amounting to 40% of the fixed adjustment cost.

Figure 23: The Recession Trap



(a) State-Dependent Effects of Shocks (b) Mechanism: Distribution & Threshold

Note: Figure 23a and 23b are based on simulations. We simulate 100,000 panels, and each panel contains 120 periods and 5,000 innovators. Boom is defined as ϱ_t equals to $-std(\varrho_t)$ at period $t = 100$; recession is defined as ϱ_t equals to $std(\varrho_t)$ at period $t = 100$. Figure 23a depicts the average distribution of $s_{i,t}$ in period $t = 100$. We re-simulated the subsamples of boom and recession with an additional 25-basis-point monetary shock imposed at period $t = 101$. Figure 23b plots the average impulse response of the boom and recession subsamples.



(a) Full Set (b) OECD Countries

Figure 24: The depth of recessions and the degree of moderation in the 2008 Recession (TFP)

Note: This 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 0.34 ($p = 0.04$) for the full sample and 0.37 ($p = 0.08$) for the OECD sample.

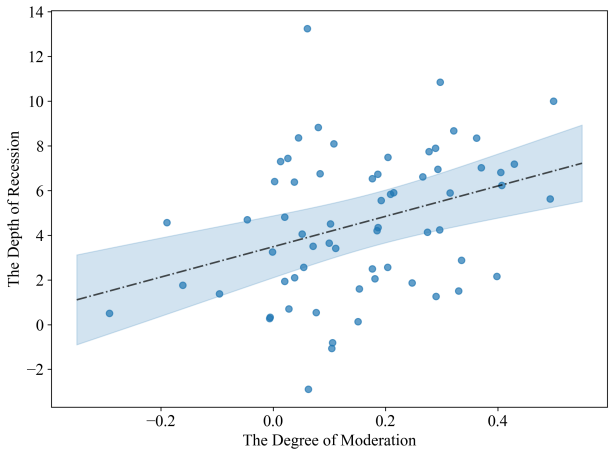


Figure 25: The depth of recessions and the degree of moderation across time

Notes: We use data in Jordà et al. (2017), which is available at the Jordà-Schularick-Taylor Macrohistory Database. The full sample set covers 18 advanced economies since 1870 on an annual basis. We focus on the major recessions of these countries after WWII. Denote t_0 the recession period. To calculate the depth of recessions, the average growth rate is the mean of the growth rate between $[t_0 - 15, t_0)$. We normalize the depth of recession by the pre-recession trend growth. The degree of moderation is the standard deviation between $[t_0 - 15, t_0)$ divided by the standard deviation between $[t_0 - 30, t_0)$. The red dotted line reports an OLS fit to the data points. The red shaded area depicts the 95% confidence band. 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 ($p = 0.00$).

B Additional Tables

Parameters	Value	Description
β	0.96	Discount rate
φ	0.5	Elasticity of unskilled labor
η	0	Elasticity of skilled labor
σ	1	Intertemporal substitution
ρ_ϱ	0.4	AR1 coefficient of the preference shock
μ	1.1	Retailers' markup on final goods
ϑ	2	Innovators' markup on technology goods
κ	90	Retailers' price adjustment cost
φ_π	1.69	Taylor rule (Inflation)
γ	0.5	Share of the technology good
ϕ_A	0.97	Depreciation rate of technology
λ	0.2	The average adoption rate of technology

Table 2: Parameters Calibrated Externally

C Data Appendix

C.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 that fall under the following criteria: not incorporated in the United States; not reported in U.S. dollars; belonging to the sectors of Utilities (SIC 4900-4999), Financial Industry (SIC 6000-6999), Non-operating Establishments (SIC 9995), or Industrial Conglomerates (SIC 9997); acquisitions exceeding 5% of total assets; firms with fewer than five R&D expense observations; firms with no positive R&D expense observations; and negative R&D expenses.

C.2 The Decomposition of the Pro-cyclicality of the Aggregate R&D

The correlation between the growth of R&D and GDP growth can be decomposed as the following:

$$\begin{aligned} Corr\Big(\frac{\Delta R\&D_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\Big) = Corr\Big(\frac{\Delta R\&D^{\text{non-Spike}}_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\Big) \frac{\sigma\big(\frac{\Delta R\&D^{\text{non-Spike}}_{t+1}}{R\&D_t}\big)}{\sigma\big(\frac{\Delta R\&D_{t+1}}{R\&D_t}\big)} \\ + Corr\Big(\frac{\Delta R\&D^{\text{Spike}}_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\Big) \frac{\sigma\big(\frac{\Delta R\&D^{\text{Spike}}_{t+1}}{R\&D_t}\big)}{\sigma\big(\frac{\Delta R\&D_{t+1}}{R\&D_t}\big)} \end{aligned} \tag{C.1}$$

where $R\&D^{\text{non-Spike}}_{t+1}, R\&D^{\text{Spike}}_{t+1}$ represents the changes of R&D due to spikes and non-spikes. Figure 3b presents the results of the decomposition analysis. Table 3 presents the detailed results of the decomposition.

Table 3: Correlation Decomposition of R&D and GDP Growth

	Aggregate R&D	Non-Spike	Spike		
			Overall	Positive	Negative
Correlation	0.51	0.43	0.44	0.36	0.19
Ratio of Std. Dev.	—	0.30	0.89	0.79	0.51
Fraction (%)	100%	24.6%	75.4%	—	—
	100%	24.6%	—	56.2%	19.2%

C.3 Lumpy R&D: Robustness Checks

Alternative Definitions of Spikes In Section 2, we identify an R&D observation as a spike if the R&D expense increases by at least 20%. Table 4 reports several alternatives. The results suggest that raising the threshold for identifying spikes does not eliminate the pro-cyclicality of the spike rate. Interestingly, the result also shows that the spike rate defined by the R&D-to-sale ratio is also positively correlated with aggregate output. This suggests that the lumpy changes in R&D intensity also exhibit a pro-cyclical pattern.

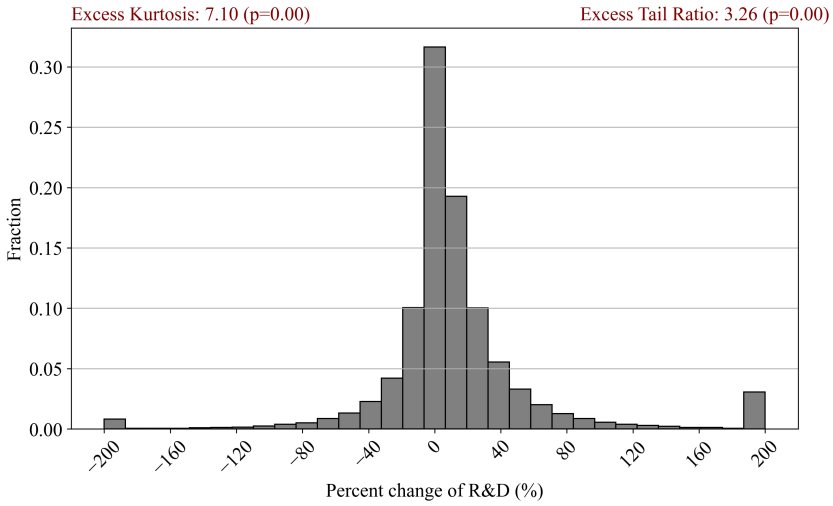
Table 4: Correlation of R&D Spike Definitions with GDP Growth

Spike Definition	Corr. GDP Growth	Corr. HP-filtered GDP	Average Spike Rate
$\Delta R\&D > 0.2$	0.57	0.66	0.24
$\Delta R\&D > 0.3$	0.60	0.67	0.18
$\Delta R\&D > 0.5$	0.56	0.64	0.11
$\Delta R\&D > 1.0$	0.45	0.45	0.06
$\Delta R\&D$-to-sale > 0.5	0.26	0.30	0.08
$\Delta R\&D$-to-sale > 1.0	0.50	0.46	0.04

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 26 and 27 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 26 demonstrates that large R&D firms also display significant lumpiness in their R&D investments. Figure 27 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.

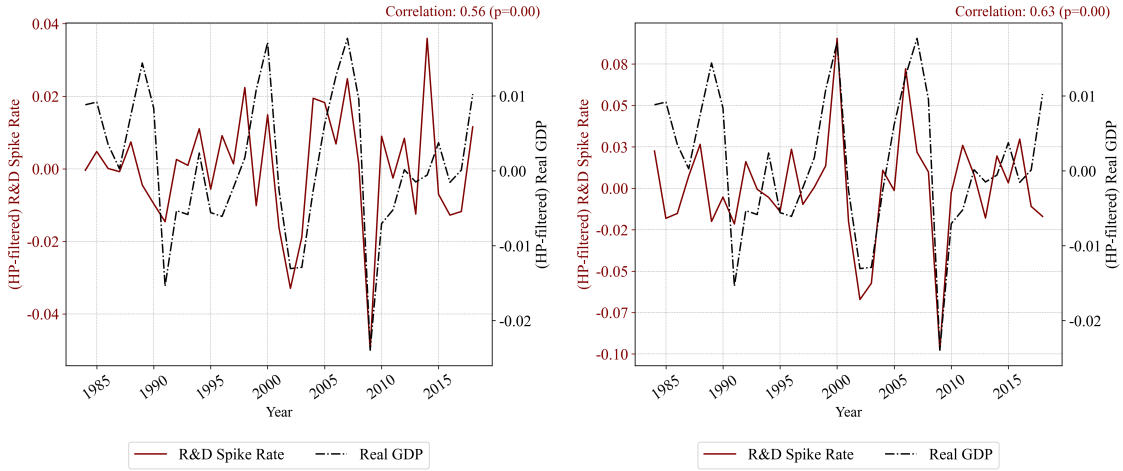
The Correlation with GDP Net of Aggregate R&D Spending Another concern related to the cyclicality of R&D investment is its inclusion as part of GDP. To verify that our findings on the cyclicality of lumpy R&D investment are not misleadingly influenced by its correlation with the R&D component of GDP,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

Figure 26: Fact 1 Robustness: 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.

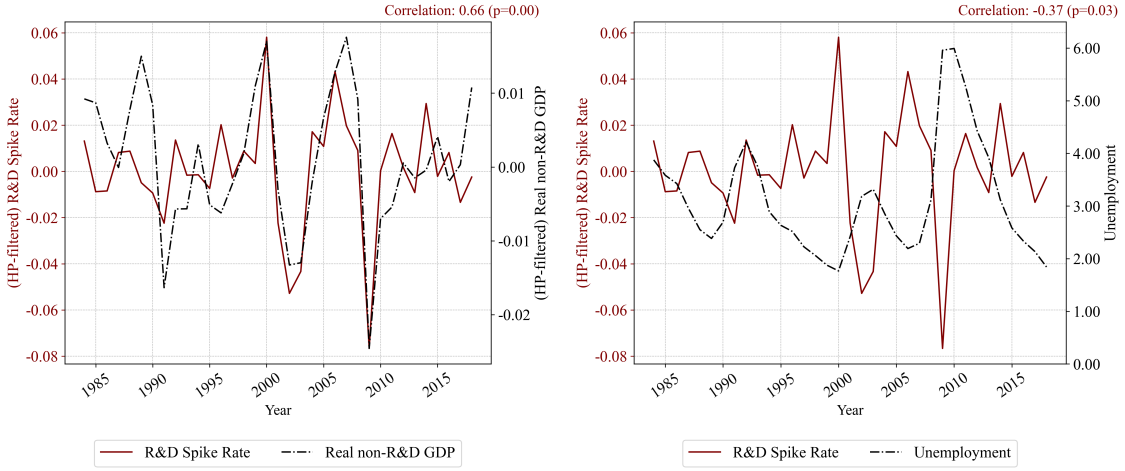
Figure 27: Fact2 Robustness: Small v.s. Large R&D Firms



(a) Small Firms (b) Large 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.

Figure 28: Fact2 Robustness: pro-cyclicality Net of Aggregate R&D Spending



(a) GDP net of R&D Spending (b) Unemployment

Note: 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.

28 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.

C.4 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:

$$\text{Gap}_{t_0,t_1} = \log Y_{t_0} + g_m \times (t_1 - t_0) - \log Y_{t_1} \tag{C.2}$$

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:

$$\text{Depth of Recession}_{t_0} = \frac{1}{K} \sum_{k=10-K+1}^{10} \text{Gap}_{t_0, t_0+k} \tag{C.3}$$

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 In Figure 11, we observe cross-country evidence suggesting a positive correlation between moderation and recession. An alternative interpretation of this evidence might be that the Great Moderation coincided with significant leveraging and a savings glut (Bagliano and Morana 2017), indirectly contributing to the recession. To test whether these financial factors could explain our findings, we propose the following regression:

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

where X_i^{2007} additionally controls for the Real Interest Rate in 2007 and the ratio of Private Credit to GDP in 2007 compared to 2000. The results are presented in Table 5. We find that including these two control variables does not alter the significance of the coefficient for moderation.

D Model Appendix

D.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

Table 5: The Robustness of the Great Twins Empirics

	(1)	(2)	(3)	(4)
Degree of Moderation	5.36***	5.01***	3.96**	4.00**
	(1.50)	(1.68)	(1.76)	(1.88)
Real interest rate$_{2007}$		0.13		0.16
		(0.26)		(0.27)
$\frac{(\text{Private Credit/GDP})_{2007}}{(\text{Private Credit/GDP})_{2000}}$			0.08	0.07
			(0.07)	(0.07)
const	3.08***	3.06***	3.29***	4.18***
	(0.68)	(0.75)	(1.19)	(1.31)
R-squared	0.24	0.26	0.15	0.21
N	42	35	35	30

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.

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 equation (3.21), which implies the intertemporal substitution of consumption. To negate the impact of the growth of A_t, we assume a monetary policy, as implied by equation (3.26), to track the efficient rate.

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. 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 a discount factor that includes a growth

component of A_t.

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.

Exogenous Inflow of Ideas In our model, idea inflow is exogenous; innovators randomly receive a certain amount of ideas each period and accumulate them until they have enough to make launching a new R&D project profitable. This simplification is designed to focus the model's main attention on lumpy R&D and its cyclicality. In fact, endogenizing idea creation would actually reinforce the mechanism described in Section 5.1. Consider pro-cyclical idea creation: during a recession, fewer ideas are generated,

leading to a reduction in the mass to the right of the adjustment threshold in Figure 21a, thereby weakening the stretching effect caused by the recession. On the other hand, during a boom, more ideas are created, resulting in a greater number of innovators initiating new R&D projects, which enhances the compressing effect of the boom.

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 Section 5.1. 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.

D.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 satisfies equation (3.5).
2. Intermediate Input Producer: Taking other prices as given, the intermediate input producer optimally sets the price that maximizes (3.7).
3. Innovator: Taking all prices as given, V_t, J_t, $\tilde{J}_t$ solves the Bellman equation (3.14), (3.13), and (3.12). $S_{i,t}$ solves the Bellman equation (3.8), 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 New Keynesian Phillips Curve (equation (3.25)) holds.
6. Central Bank: The Taylor rule (equation (3.26)) holds.

7. The distribution of $s_{i,t}$ and $x_{i,t,j}$ evolves according to the adjustment threshold $\bar{s}_t$.
8. A_t evolves according to equation (3.18).
9. All markets (final good, intermediate good, intermediate input, labor, skilled labor) clear.

Linearization Based on the discussion in the Online Appendix D.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:

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

Household's consumption Euler equation:

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

The Intermediate goods producer's demand for M_t:

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

The Intermediate goods producer's demand for L_t:

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

Retailer's price-setting problem (NKPC):

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

Monetary policy:

$$i_t = r_t^e + \varphi_\pi \pi_t \tag{D.6}$$

Efficient real rate of interest[10]:

$$r_t^e = E_t g_{t+1} - (1-\omega) E_t \Delta \varrho_{t+1} \tag{D.7}$$

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:

$$\tilde{c}_t = \log(1-\gamma) + \tilde{y}_{t,m} \tag{D.8}$$

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 (3.14), (3.13), (3.12), (3.9), (3.10), and (3.8), 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.

D.3 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

$$S(s) = I\{s < \bar{s}\}(\beta E[S(\phi s + \eta)]) + (1 - I\{s < \bar{s}\})(-c + sJ^A + \beta E[S(\eta)]) \tag{D.9}$$

Next we show that equation (D.9) is preserved by the recursive representation.

Denoting

$$f(s) = -c + sJ^A + \beta E[S(\eta)], \quad g(s) = \beta E[S(\phi s + \eta)] \tag{D.10}$$

[10]Without loss of generality, we set $\omega = 1$ in solving and calibrating the model. Because it is always equivalent to redefining the preference shock under the efficient-rate Taylor rule as $\hat{\varrho}_t \equiv \omega \varrho_t$.

For $f(s)$,

$$f'(s) = J^A \tag{D.11}$$

for all $s > 0$.

We rewrite $g(s)$ as

$$\begin{aligned} g(s) =& Pr(\phi s + \eta > \bar{s})\beta[-c + (\phi s + E[\eta | \phi s + \eta > \bar{s}]) + \beta E[S(\eta)]] \\ &+ Pr(\phi s + \eta \leq \bar{s}, \phi^2 s + \phi\eta + \eta_1 > \bar{s})\beta^2[-c + (\phi^2 s + E[\phi\eta + \eta_1 | \phi^2 s + \phi\eta + \eta_1 > \bar{s}]) + \beta E[S(\eta)]] \\ &+ ... \end{aligned}$$

which can be rewritten as

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

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 $s \geq 0$. 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.4 The Proof of Proposition 3

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

$$s_{i,t+1} = \phi s_{i,t} I\{s_{i,t} < \bar{s}\} + \eta_{i,t+1}, \tag{D.12}$$

The implied equation for the steady-state distribution of $s_{i,t}$ is given by:

$$F(s;\bar{s}) = \begin{cases} \tilde{P}_{\bar{s}}\Psi(s) + \int_0^{\frac{1}{\phi}s} f(x;\bar{s})\Psi(s-\phi x)dx, & \text{if } s \leq \phi\bar{s} \\ \tilde{P}_{\bar{s}}\Psi(s) + \int_0^{\bar{s}} f(x;\bar{s})\Psi(s-\phi x)dx, & \text{otherwise} \end{cases} \tag{D.13}$$

where $\tilde{P} \equiv \int_{\bar{s}}^{\infty} f(s;\bar{s})ds$.

We assume the density of idea inflow, $\psi(\eta)$, are well-behaved, that $\psi \in C^1([0,\infty))$. To keep the discussion simple, we further assume ψ is unimodal. The proof for multi-modal ψ is similar, but one has to focus on the f in the region beyond the second right-most modal of ψ.

Equation (D.13) suggests that for $s \geq \bar{s}$,

$$f(s;\bar{s}) = \tilde{P}_{\bar{s}}\psi(s) + \int_0^{\bar{s}} f(x;\bar{s})\psi(s-\phi x)dx \tag{D.14}$$

$$f'(s;\bar{s}) = \tilde{P}_{\bar{s}}\psi'(s) + \int_0^{\bar{s}} f(x;\bar{s})\psi'(s-\phi x)dx \tag{D.15}$$

By Proposition 1, there is a one-to-one mapping between the adjustment cost c and the adjustment threshold $\bar{s}$. The proposition is equivalent to showing that, for any ψ, there is some $\bar{s}^* \geq 0$, such that $f'(s;\bar{s}) \leq 0$ for $s \geq \bar{s}$ with $f'(s;\bar{s}) < 0$ for some non-zero subset holds[11] if and only if $\bar{s} > \bar{s}^*$.

The following lemma establishes that, if $f'(\bar{s};\bar{s}) < 0$, it implies $\bar{s}$ is on the declining region of the distribution.

Lemma D.1. *If $f'(\bar{s};\bar{s}) < 0$, it follows $f'(s;\bar{s}) < 0$ for $s > \bar{s}$.*

The next lemma establishes that there is a minimum level of adjustment threshold, such that all actual adjustment thresholds beyond that level will posit on the declining region of the idea distribution.

Lemma D.2. *There is a unique $\bar{s}^* > 0$, such that*

$$f'(\bar{s}^*;\bar{s}^*) = \tilde{P}_{\bar{s}^*}\psi'(\bar{s}^*) + \int_0^{\bar{s}^*} f(x;\bar{s}^*)\psi'(\bar{s}^*-\phi x)dx = 0 \tag{D.16}$$

and for $\bar{s} > \bar{s}^*$, $f'(\bar{s};\bar{s}) < 0$.

[11] In other words, $\bar{s}$ is on the right of the last modal of f.

By Lemma D.1 and D.2, if $\bar{s} > \bar{s}^*$, $f'(s;\bar{s}) < 0$ for $s > \bar{s}$. On the other hand, if $\bar{s} < \bar{s}^*$, it follows $f'(\bar{s};\bar{s}) > 0$, and due to the continuity, there must be some $s > \bar{s}$ that $f'(s;\bar{s}) > 0$. This suggests $f'(s;\bar{s}) < 0$ for $s > \bar{s}$ holds if and only if $\bar{s} > \bar{s}^*$. This completes the proof of Proposition 3.

Next, consider the case where $\psi' < 0$. This condition suggests that for any $\bar{s} > 0$:

$$f'(s;\bar{s}) = \underbrace{\tilde{P}_{\bar{s}}\psi'(s)}_{<0} + \int_0^{\bar{s}} \underbrace{f(x;\bar{s})}_{>0}\underbrace{\psi'(s-\phi x)}_{<0}dx < 0 \tag{D.17}$$

holds for all $s \geq \bar{s}$. This suggests as long as $c > 0$ and $\psi' < 0$, the condition for the asymmetric effect of shocks on the distribution holds, which completes the proof of Corollary 1.

D.5 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:

$$\text{ASY}_\varepsilon(\bar{s}) \equiv \frac{\int_{\bar{s}(1-\varepsilon)}^{\bar{s}} f(x)dx - \int_{\bar{s}}^{\bar{s}(1+\varepsilon)} f(x)dx}{\int_{\bar{s}(1-\varepsilon)}^{\bar{s}(1+\varepsilon)} f(x)dx} \tag{D.18}$$

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.13), the steady state distribution of $s_{i,t}$ solves:

$$f(s) = \tilde{P}\exp\{-s\} + \int_0^s f(s)\exp\{-s+x\}dx \tag{D.19}$$

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 $\int_0^{\bar{s}} f(s)dx = \tilde{P}\bar{s}$:

$$\tilde{P} = \frac{1}{\bar{s}+1} \tag{D.20}$$

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

$$\text{ASY}_\varepsilon(\bar{s}) = \frac{\bar{s} - 1 + \exp\{-\bar{s}\varepsilon\}}{\bar{s} + 1 - \exp\{-\bar{s}\varepsilon\}} \tag{D.21}$$

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:

$$\psi(\eta) = \frac{1}{2}(1+x)\exp(-x) \tag{D.22}$$

Using equation (D.13), the steady state distribution of $s_{i,t}$ solves:

$$f(s) = \frac{1}{2}\tilde{P}(1+s)\exp(-s) + \int_0^s \frac{1}{2} f(s)(1+s-x)\exp\{-s+x\}dx \tag{D.23}$$

One can show that the solution:

$$f(x) = \begin{cases} \tilde{P}\left(\frac{4}{3} - \frac{1}{6}\exp\left(-\frac{3}{2}x\right)\right) & \text{for } s \leq \bar{s}, \\ \tilde{P}\left(\frac{4}{3}\exp(-\bar{s}) - \frac{1}{6}\exp\left(-\frac{1}{2}\bar{s}\right)\right)\exp(-s) & \text{for } s > \bar{s}. \end{cases} \tag{D.24}$$

Then it follows:

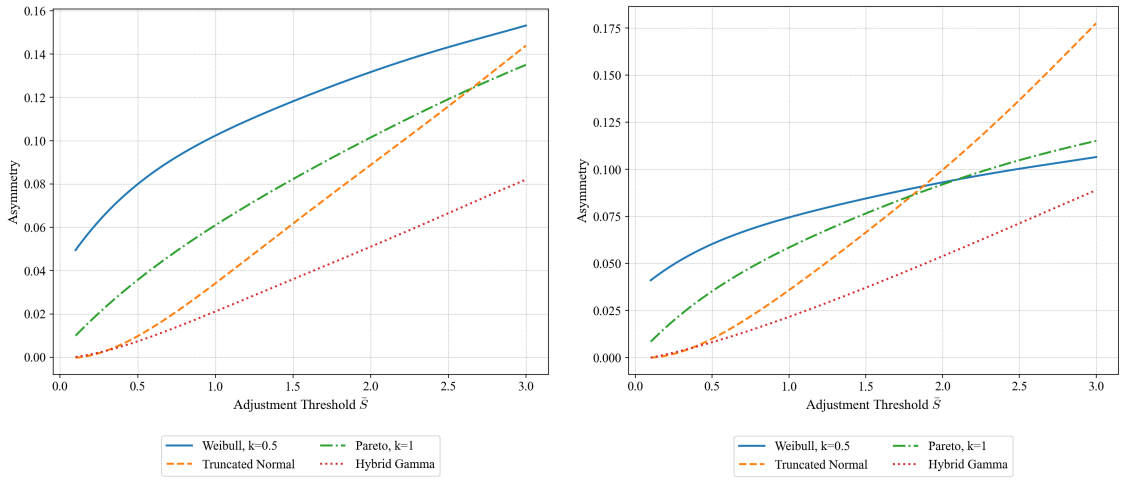
$$\int_{\bar{s}}^{\bar{s}(1+\varepsilon)} f(x)dx = \tilde{P}\Big[\frac{4}{3}(1-\exp(-\bar{s}\varepsilon)) - \frac{1}{6}\exp(-\frac{3}{2}\bar{s})\big(1-\exp(-\bar{s}\varepsilon)\big)\Big] \tag{D.25}$$

$$\int_{\bar{s}(1-\varepsilon)}^{\bar{s}} f(x)dx = \tilde{P}\Big[\frac{4}{3}\bar{s}\varepsilon - \frac{1}{6}\exp(-\frac{3}{2}\bar{s})(\frac{2}{3}\exp(\frac{3}{2}\bar{s}\varepsilon) - \frac{2}{3})\Big] \tag{D.26}$$

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

$$\text{ASY}_\varepsilon(\bar{s}) = \frac{\frac{2}{3} - \frac{7}{12}\exp(-\frac{3}{2}\bar{s})}{\frac{8}{3} - \frac{1}{3}\exp(-\frac{3}{2}\bar{s})} \cdot (\bar{s}\varepsilon) \tag{D.27}$$

Figure 29: 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 29 demonstrates that the positive relationship between asymmetry and lumpiness persists across different model specifications.

D.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:

$$y_{i,t} = (Z_t e_{i,t})^{1-\alpha} k_{i,t}^{\alpha}, \quad \text{where} \quad \alpha < 1. \tag{D.28}$$

Productivity components evolve according to:

$$z_{t+1} = \varepsilon^z_{t+1}, \quad \log e_{i,t+1} = \log e_{i,t} + \varepsilon^e_{i,t+1}. \tag{D.29}$$

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:

$$\theta_{i,t} = Z_t e_{i,t} \Theta(i_t, \Delta N_t). \tag{D.30}$$

where N_t is a Poisson counter with arrival rate λ:

$$\Theta(i_t, \Delta N_t) = \begin{cases} 0 & \text{if } i_t = 0 \text{ or } \Delta N_t = 1, \\ \theta^- & \text{if } i_t > 0 \text{ and } \Delta N_t = 0, \\ \theta^+ & \text{if } i_t < 0 \text{ and } \Delta N_t = 0. \end{cases} \tag{D.31}$$

Assuming no capital depreciation, capital stock evolves according to:

$$k_{i,t} = k_{i,t-1} + i_{i,t}. \tag{D.32}$$

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

$$V(k_{i,t-1}; Z_t, e_{i,t}) = \max_{i_{i,t}} \left\{ (Z_t e_{i,t})^{1-\alpha} k^\alpha_{i,t} - r k_{i,t} - \theta_{i,t} + \beta E_t[V(k_{i,t}; Z_{t+1}, e_{i,t+1})] \right\}. \tag{D.33}$$

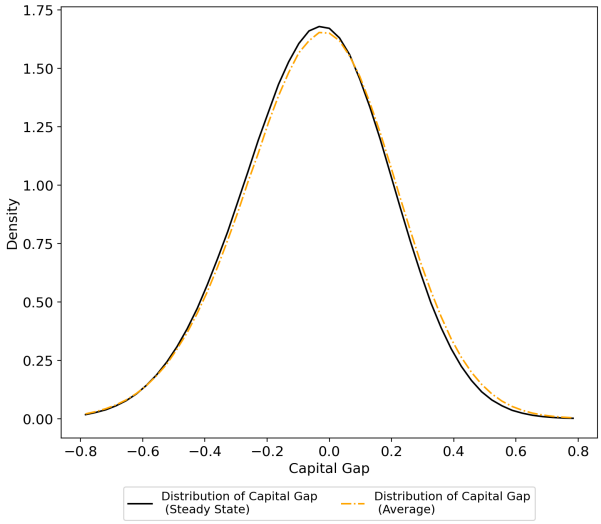
The investment policy features an inactive region $\mathcal{R}$ defined by $\left(k^-(Z,e), k^+(Z,e)\right)$ and an optimal capital level $k^*(Z,e)$, where no capital adjustment occurs within $\mathcal{R}$; outside this region:

$$i_{i,t} = k^*(Z_t, e_t) - k_{i,t-1}. \tag{D.34}$$

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

Figure 30: Steady State v.s. Average: Calibrated



Note: We calibrate the model to match the (annual) distribution of capital investment in Compustat since 1983. The model is calibrated quarterly. The resulting parameter values are: the capital depreciation $\delta = 0.077/4$, $\sigma_e = 0.2$, $\log(x^*) - \log(x^+) = -0.4$, $\log(x^*) - \log(x^-) = 0.15$, $\lambda = 0.01$. We allow for persistence in the aggregate shock of z_t, where $\rho_z = 0.95$ and $\sigma_z = 0.03$.

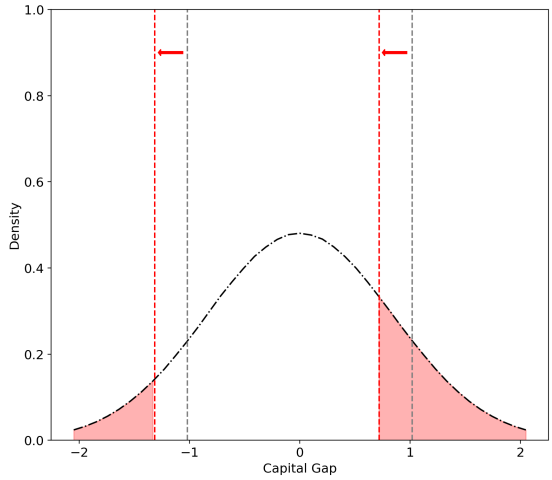
ideas in our lumpy R&D model, where the capital gap's dynamics are represented as:

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

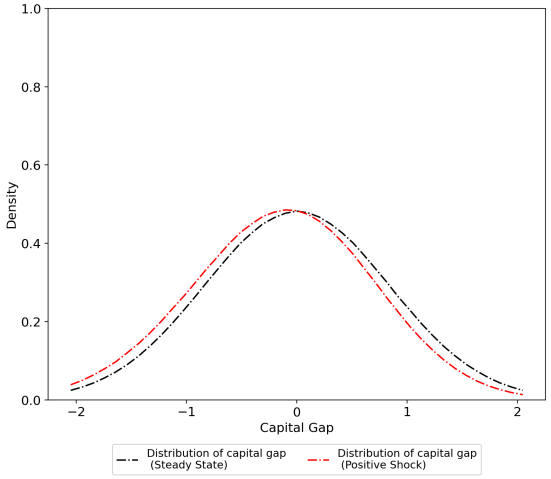
In a more realistic scenario, factors such as capital depreciation, partial irreversibility, and random opportunities for free adjustments may also influence the impact of aggregate shocks on the distribution of the capital gap. Therefore, in Figure 30, we calibrate the parsimonious lumpy investment 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 30 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.

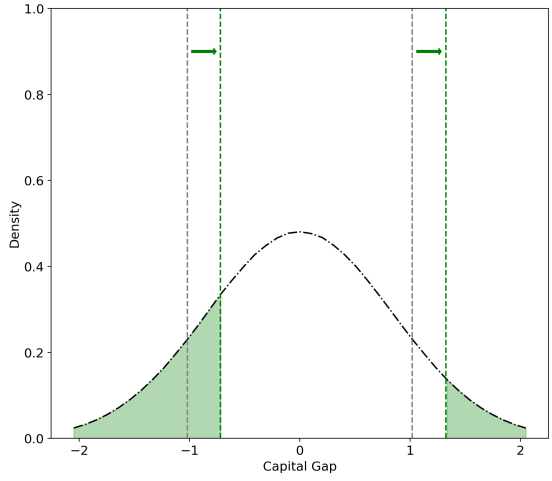
Figure 31: The Effect of Aggregate Shocks on the Distribution of Capital Gap



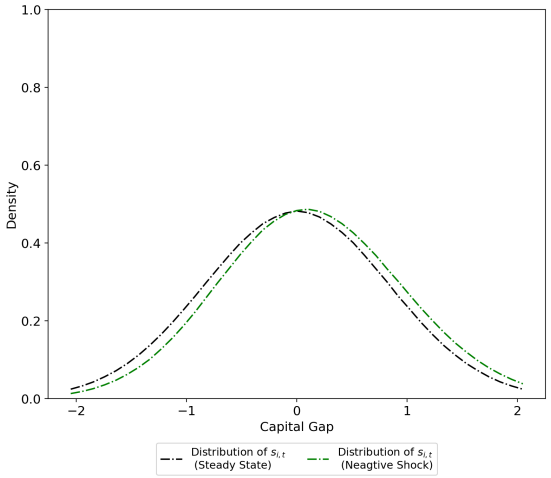
(a) Positive shocks shift inactive region to the left



(b) The distribution after a positive shock



(c) Negative shocks shift inactive region to the right



(d) The distribution after a negative shock