LUMPY INNOVATION AND THE GREAT TWINS: GREAT MODERATION AND GREAT RECESSION

LIXING WANG
University of Bonn and National University of Singapore

DONGHAI ZHANG
National University of Singapore

We document that firm-level innovation decisions 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 innovating firms and fixed adjustment costs, featuring lumpy innovation. Through this framework, we uncover a novel mechanism underpinning the 'Great Twins' hypothesis: the Great Moderation–marked by diminished economic volatility up to 2008–causally intensified the severity of the subsequent Great Recession. We provide micro- and macro-level evidence supporting this hypothesis. Finally, we explore the policy implications, including a new tail-risk and variance tradeoff for central banks.

KEYWORDS: Heterogeneous Innovators, Lumpy Innovation, Great Twins, Medium-run Business.

Lixing Wang: lixing.wang@uni-bonn.de

Donghai Zhang: donghai.d.zhang@gmail.com

We thank Diego Comin, Ryan Chahrour, John Fernald, Pablo A. Guerrón-Quintana, Franck Portier, Yongseok Shin, Ippei Fujiwara, Michael Song, Vincent Sterk, Daniel Xu, as well as the seminar/conference participants at NBER EASE 2024, Collegio Carlo Alberto, Bank of Portugal, University of Bonn, Singapore Management University, National University of Singapore, Fudan University, SUFE, Jinan University (IESR), Academia Sinica, and Virtual East Asian Macroeconomics Seminar, Southern Workshop in Macroeconomics for their suggestions and discussions. Donghai Zhang thanks the Singapore Ministry of Education for financial support under the SSRTG grant SSRC2023-SSRTG-026. Lixing Wang thanks the DFG for financial support as part of the research training group RTG 2281 'The Macroeconomics of Inequality'.

1. INTRODUCTION

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

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

To formally motivate our theoretical framework, we document three novel facts. First, the distribution of changes in at the firm level is heavily populated by zeros/small values and large values—a lumpy pattern. Second, the fraction of firms making big adjustments (or spike adjustments) in R&D expenses—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 innovation. The model can replicate the firm-level distribution of changes in R&D expenses 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—causally contributed to 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 raw ideas they have. In the model, the amount of ideas is isomorphic to the quality of an idea. The function of an R&D project is to process the stock of raw ideas into R&D ideas. We distinguish between raw ideas and R&D ideas, where the latter refers to ideas actively pursued within an on-going R&D project. To create a new R&D project (to innovate), the innovating firm incurs a fixed cost. The key decision innovators make during the innovation phase is whether to set up a new R&D project, i.e., the extensive margin innovation decision. To convert R&D ideas into new technologies, the source of growth, innovators need to hire scientists. Scientists' wage bills constitute the variable cost. This describes the second phase, called the adoption phase. In the last phase, the monetizing phase, innovators claim the intellectual property right (patent) of the new technologies and sell the patent to an intermediate input producer.

The model incorporates standard business cycle features that permit the analysis of the interaction between business cycle fluctuations and R&D decisions. Business cycle fluctu-

ations affect the profits of new technologies and, therefore, alter innovators' decisions, particularly along the extensive margin. The model predicts that R&D spike adjustments are pro-cyclical, consistent with our empirical observations. The model is calibrated to match key moments in the U.S. economy and is shown to match untargeted moments, including the motivating facts described above, and the distribution of innovators. Furthermore, we fit the model to the evolution of real GDP 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 *qualitative* insight: the Great Moderation that pre-dated the Great Recession has contributed to the severity of the latter — a phenomenon labeled as the *Great Twins*. Specifically, we show that a decrease in the sample volatility of GDP in the years preceding the Great Recession, either due to Good Luck (drawing smaller shocks) or Good Policy (U.S. Federal Reserve adopted a more effective monetary policy to reduce economic fluctuations), *causally* increases the effects of shocks that triggered the Great Recession on productivity and real GDP. The Great Twins hypothesis has an important implication for policymakers: a 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 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. Lastly, using citations as a proxy for patent quality, we show that the average patent quality has increased since the Great Recession, consistent with an natural outcome of the extensive margin innovation decisions we highlight: innovations conducted during a recession have, on average, higher quality.

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).[1] 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

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

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.

There is a long tradition of analyzing the aggregate implications of firms' lumpy capital investment behavior with fixed adjustment costs; see, for example, Caballero et al. (1995), Caballero and Engel (1999), Khan and Thomas (2008), Bachmann et al. (2013), Miao and Wang (2014) and more recently, Koby and Wolf (2020), Winberry (2021), Baley and Blanco (2021), Gnewuch and Zhang (2023), Chiavari and Goraya (2024) and Lippi and Oskolkov (2023). To the best of our knowledge, the new mechanism in our model is unique within the context of lumpy innovation and has not been previously documented in the literature.[2] 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), Barlevy (2007), Broda and Weinstein (2010), Ouyang (2011), and Fabrizio and Tsolmon (2014) document that aggregate R&D spending is pro-cyclical, with this pro-cyclicality being observed across different contexts and for various sources

[2]Unlike the typical lumpy capital investment (Ss) model, our lumpy innovation exhibits a one-S feature, which is crucial to the mechanism behind the Great Twins: the asymmetric effects of shocks on the distribution. Secondly, since the incentive for innovation primarily comes from innovation profits, our model does not encounter the issue of irrelevance of fixed costs due to the semi-elasticity of investment with respect to the real interest rate approaching infinity, as seen in the capital investment literature (see, e.g., Caballero and Engel (1999), Khan and Thomas (2008), Koby and Wolf (2020), Winberry (2021)).

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 nonlinear mechanism—the causal effects of reduced volatility (a sequence of small shocks) on increased left-tail risk—in our model is an *endogenous* feature. Moreover, different from the existing literature, we emphasize R&D investment along the extensive margin, consistent with our empirical findings that the pro-cyclicality of R&D investments is primarily due to the extensive margin. Overall, we view the main contribution of our paper as providing a new qualitative insight.

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 changes in R&D investments of innovating 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 innovating firms. In the empirical analysis, a firm is classified as an innovating firm if it reports positive R&D expenses at least once in our sample period. For cross-time analysis, we retain firms with a minimum of five observations. We also exclude acquisitions exceeding 5% of total assets.[3] Our sample consists of 9,166 innovating firms, representing 42.8% of the companies in the Compustat universe. The industrial composition of our sample is similar to the full Compustat sample, as detailed in Table 1 of Online Appendix. The panel is unbalanced, with each firm averaging 16.6 observations.

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

Formalizing the Concept of Lumpy Innovation To formalize the concept of lumpy innovation, Online Appendix presents 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, innovating firms make a spike adjustment in R&D. We refer to these observations as evidence of lumpy innovation.

An innovation, or starting a new R&D project, requires a significant change in R&D expenditure both today and in the future due to, for example, the hiring of new scientists and their associated wage bills. Formally, lumpy innovation postulates that the distribution

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

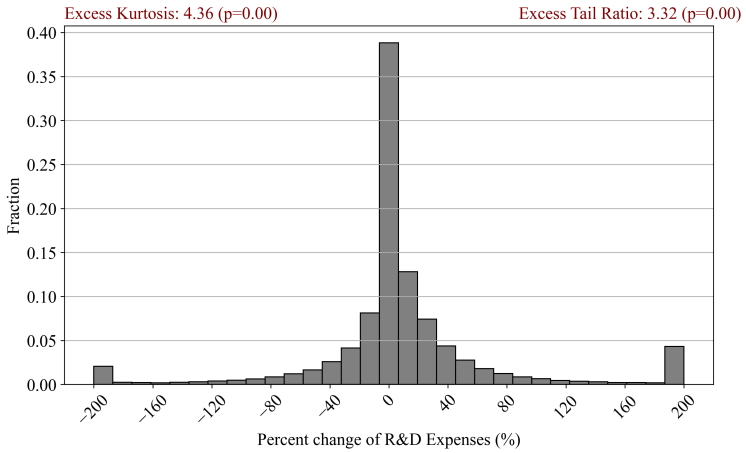
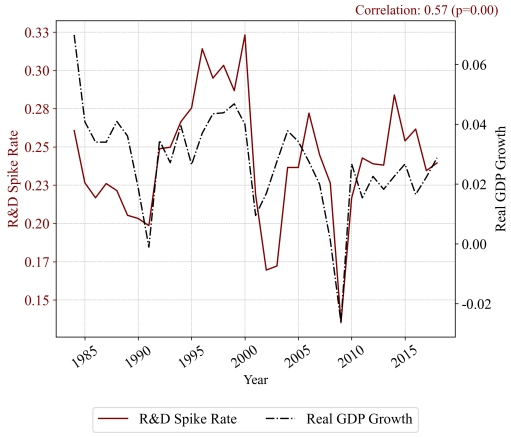


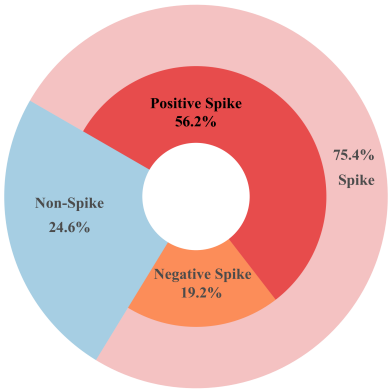
FIGURE 1.—Firm-Level Distribution of Percent Change of R&D Expenses. Note: To handle cases with zero initial R&D expenses, we use the formula: percent change $= \frac{\text{R\&D}_t - \text{R\&D}_{t-1}}{0.5 \times (\text{R\&D}_t + \text{R\&D}_{t-1})}$. Excess kurtosis measures deviation from a normal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

of *changes* in R&D is heavily populated by zeros/small values and large values. Is lumpy innovation a general phenomenon in the data? This question is addressed in the following section.

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 innovation: R&D investment 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. In Online Appendix 1.3, among other robustness checks, we discuss that the lumpy pattern is not due to strategic accounting for tax reasons. In fact, U.S. R&D Tax Credit rules discourage it. Moreover, the lumpy pattern is not due to entry/exit: the excess kurtosis is 6.06 ($p = 0.00$), and the excess tail ratio is 1.20 ($p = 0.00$) if entry/exit observations are excluded.





(a) R&D Spike Rate and Real GDP growth (b) Decompose the pro-cyclicality

FIGURE 2.—The Pro-cyclicality of R&D spikes. Note: Panel (a) plots the time series of the R&D spike rate, defined as the fraction of R&D spikes and real GDP from 1984 to 2018. The decomposition in Panel (b) is based on the simple observation that changes in R&D expenses can be decomposed into spike and non-spike changes, such that $Cov\left(\frac{\Delta R\&D_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right) = Cov\left(\frac{\Delta R\&D^{Spike}_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right) + Cov\left(\frac{\Delta R\&D^{non-Spike}_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right)$. See Online Appendix for the details of the decomposition.

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

2.2. *The Cyclicality of Lumpy Innovation*

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 ex-

penses.[4] The R&D spike *rate* is defined as the fraction of innovating 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 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, 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. In Online Appendix 1.3, we show 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, the aggregate R&D spending is strongly pro-cyclical, with a correlation of 0.51 (p-value $= 0.00$)[5].

Is the lumpy nature of R&D investment at the micro-level relevant in explaining the cyclicality of aggregate R&D expenditures? To address this question, we 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, such that $Cov\left(\frac{\Delta R\&D_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right) = Cov\left(\frac{\Delta R\&D^{Spike}_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right) + Cov\left(\frac{\Delta R\&D^{non-Spike}_{t+1}}{R\&D_t}, \frac{\Delta Y_{t+1}}{Y_t}\right)$, where $R\&D^{\text{non-Spike}}_{t+1}$ and $R\&D^{\text{Spike}}_{t+1}$ represent the changes in R&D due to non-spike and spike adjustments, respectively. See Online Appendix 1.3 for the details of the decomposition.

Figure 2b 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

[4] 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 2b shows that negative spikes contribute less to R&D's pro-cyclicality compared to positive spikes.

[5] In Online Appendix, 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).

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 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.

Robustness Checks In Online Appendix 1.3, we provide a detailed discussion on additional robustness checks for Facts 1-3. Firstly, we provide evidence that the lumpy pattern observed in Figure 1 is not merely a reflection of volatile changes in firms' sales, and the lumpy pattern also holds for patent data. Secondly, we show that our results are robust across various alternative definitions of R&D spikes. Thirdly, 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. Fourthly, 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. Lastly, the lumpy changes in R&D are not driven by financial frictions. Specifically, we demonstrate that the lumpy pattern is qualitatively consistent and quantitatively nearly identical across groups of firms with both high and low leverage.

In the remainder of the paper, we provide a model that interprets the spike adjustment in R&D as evidence suggesting that the innovator is setting up a new innovation project—R&D investment along the extensive margin. 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.

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 innovating firms, we introduce *heterogeneous* innovating 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.

Our model is 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 $M_t(i)$:

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

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 $M_{t,j}$:

$$M_{t,j} = \left(\frac{P_{t,j}^M}{P_t^M} \right)^{\frac{\vartheta}{1-\vartheta}} M_t, \tag{2}$$

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:

$$Y_{t,m} = L_t^{1-\gamma} M_t^{\gamma}, \tag{3}$$

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

$$\max_{M_t, L_t} P_{t,m} Y_{t,m} - W_t L_t - P_t^M M_t, \tag{4}$$

subject to production function (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:

$$\gamma \left(\frac{L_t}{M_t}\right)^{1-\gamma} = \frac{P_t^M}{P_{t,m}}, \qquad (1-\gamma)\left(\frac{M_t}{L_t}\right)^{\gamma} = \frac{W_t}{P_{t,m}}. \tag{5}$$

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:

$$M_{t,j} = f(Y_{t,m}) = Y_{t,m}. \tag{6}$$

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

$$\Phi_{t,j} = (P_{t,j}^M - P_{t,m})\left(\frac{P_{t,j}^M}{P_t^M}\right)^{\frac{\vartheta}{1-\vartheta}} M_t. \tag{7}$$

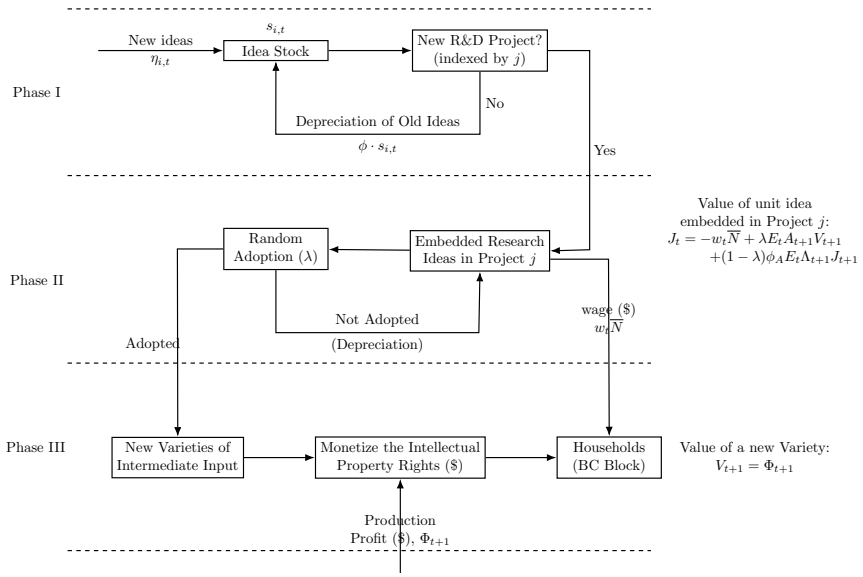
Since the production process is symmetric across intermediate input producers, $\Phi_t = \Phi_{t,j}$ for all j. Hereafter, we drop the j index for intermediate input producers and work with Φ_t, denoting the profit of any intermediate input.

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 3 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.

FIGURE 3.—Block Diagram of the Model System – the Innovation Block



This phase is characterized by the hiring of skilled labor for idea adoption, 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, or quality of raw ideas, in preparation for adoption.

There is a continuum of heterogeneous innovators, each receiving an exogenous amount (or quality improvement) $\eta_{i,t}$ 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:

$$S_t(s_{i,t}) = \max\left\{S_t^a(s_{i,t}), S_t^n(s_{i,t})\right\}, \tag{8}$$

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

$$S_t^a(s_{i,t}) = -cA_t + s_{i,t}\tilde{J}_t + E_t\Lambda_{t,t+1}S_{t+1}(\eta_{i,t+1}). \tag{9}$$

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:

$$S_t^n(s_{i,t}) = E_t\Lambda_{t,t+1}S_{t+1}(\phi s_{i,t} + \eta_{i,t+1}). \tag{10}$$

Solving innovators' problem, as described in Equations (8) - (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 Online Appendix 3.1. *Q.E.D.*

FIGURE 4.—R&D Decisions Along the Extensive Margin

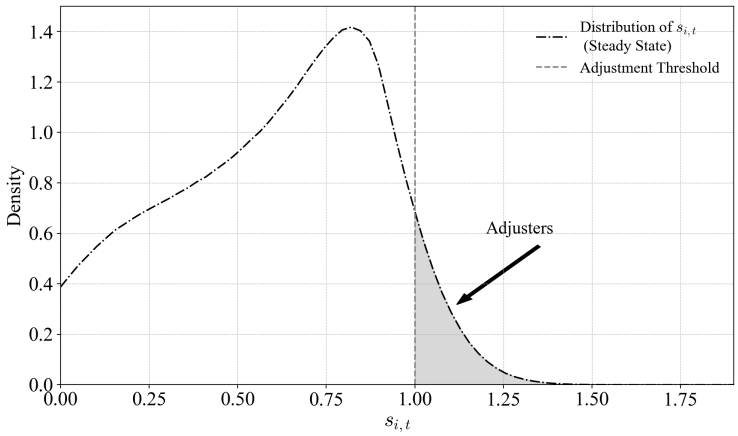


Figure 4 visualizes Proposition 1. It plots the distribution of $s_{i,t}$ together with the threshold threshold $\bar{s}_t$. The shaded area indicates the aggregate amount of adjusters—innovators who decide to set up new projects—in the economy.

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:

$$s_{i,t+1} = \phi s_{i,t} I\{s_{i,t} < \bar{s}_{i,t}\} + \eta_{i,t+1}, \tag{11}$$

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), $\Lambda_{t,t+1}$ represents 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:

$$\tilde{J}_t = \tilde{\lambda} J_t + \phi_A (1 - \tilde{\lambda}) E_t \Lambda_{t,t+1} J_{t+1}. \tag{12}$$

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 $\bar{n}$ units 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:

$$J_t = -W_t^{RD} \bar{n} + \lambda E_t \Lambda_{t,t+1} A_{t+1} V_{t+1} + (1 - \lambda) \phi_A E_t \Lambda_{t,t+1} J_{t+1}, \tag{13}$$

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 (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 an R&D project transforms into a one-period patent (Acemoglu et al. 2012), allowing the innovator to sell the patent to an intermediate input producer and claim the profit of production for a single period. The value of the patent is defined as

$$V_t = \Phi_t, \tag{14}$$

where Φ_t is defined in Equation (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 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 $\tilde{x}_{i,t}^{new}$, is:

$$\tilde{x}_{i,t}^{new} = s_{i,t} I\{s_{i,t} \geq \bar{s}_{i,t}\}. \tag{15}$$

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:

$$\tilde{X}_{t+1} = \phi_A(1-\tilde{\lambda})\tilde{X}_t + \int_i \tilde{x}_{i,t+1}^{new}\, di, \tag{16}$$

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 $\int_i \tilde{x}_{i,t+1}^{new}\, di$ 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:

$$X_{t+1} = \phi_A(1-\lambda)X_t + \tilde{\lambda}\tilde{X}_{t+1}. \tag{17}$$

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

$$A_{t+1} = \phi_A A_t + \lambda X_t A_{t+1}. \tag{18}$$

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 (15). This, in turn, ultimately affects the aggregate productivity A_t through Equations (16) to (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

$$\max_{\{C_{t+i},L_{t+i},N_{t+i},Q_{t+i}\}} E_t \sum_{i=0}^{\infty} \beta^i \varrho_{t+i} \left[\frac{C_{t+i}^{1-\sigma}}{1-\sigma} - \frac{L_{t+i}^{1+\varphi}}{1+\varphi} - \left(\frac{N_{t+i}^{1+\eta}}{1+\eta} \right) \right] \tag{19}$$

subject to

$$P_t C_t = W_t L_t + W_t^{RD} N_t + Q_t - R_t^{-1} Q_{t+1} + \Xi_t \tag{20}$$

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

$$1 = E_t \left[\Lambda_{t,t+1} \Pi_{t+1}^{-1} R_t \right] \tag{21}$$

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

$$C_t^\sigma L_t^\varphi = \frac{W_t}{P_t}, \qquad C_t^\sigma N_t^\eta = \frac{W_t^{RD}}{P_t}. \tag{22}$$

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:

$$Y_t = \left(\int y_{t,h}^{\mu} \right)^{\frac{1}{\mu}} \tag{23}$$

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

$$y_{t,h} = \left(\frac{p_{t,h}}{P_t} \right)^{\frac{\mu}{1-\mu}} Y_t. \tag{24}$$

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:

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

subject to the demand constraint (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):

$$\pi_t = \beta E_t \pi_{t+1} + \frac{\mu}{\kappa(1-\mu)} \left(\frac{1}{\mu} - \frac{P_{t,m}}{P_t} \right). \tag{25}$$

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):

$$R_t = R_t^e \Pi_t^{\varphi_\pi}, \tag{26}$$

In our framework, this efficient real interest rate emerges under the assumption of a competitive retail market[6]. The efficient real interest rate, R_t^e, is then implied by the equation of optimal intertemporal substitution:

$$1 = \beta E_t \Big[\Big(\frac{\varrho_{t+1}}{\varrho_t} \Big)^{\omega - 1} \frac{A_t}{A_{t+1}} R_t^e \Big] \tag{27}$$

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

Discussions of Model Assumptions Online Appendix 2.1 provides detailed discussions of model assumptions, including, but not limited to, the assumptions we made to make the model tractable. One aspect worth emphasizing is that, in the model, distribution matters for the evolution of aggregate productivity A_t, and therefore, distribution can affect households' consumption decisions. In this case, one would need to keep track of the distribution of $s_{i,t}$ and $x_{i,t}$. Since distributions are infinite dimensional variables, employing numerical methods to handle them would significantly increase the model's complexity and reduce its interpretability.[7] The New Keynesian block, together with the assumed Taylor rule, ensures that the real interest rate does not depend on the growth rate of the economy, and therefore, R&D decisions and distributions do not affect households' consumption decisions. This feature greatly simplifies the model and maintains the focus of our analysis on how the distribution alters R&D decisions.

Online Appendix 2.1 presents additional discussions on the model assumptions regarding the depreciation of ideas, the exogenous inflow of ideas, passive adaptation, and the time-to-build.

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

[7] 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.

3.4. *Market Clearing*

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

$$Y_t = C_t, \qquad Y_{t,m} = \int_0^1 y_{t,h} dh + \int_0^{A_t} M_{t,j} dj. \tag{28}$$

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 $Q_t = 0$. Online Appendix 2.2 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.

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

Next, we turn to the parameter that governs the obsolescence of unused raw ideas, $1-\phi$. It determines the difficulty of further accumulating raw ideas as the idea stock approaches the adjustment threshold. The larger the obsolescence rate, the less likely that an R&D project contains an amount of raw ideas much larger than the adjustment threshold $\bar{s}_t$. In other words, a larger $1-\phi$ will lead to a smaller spike in R&D expenses. Therefore, we include the 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 I. 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 I.

TABLE I

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

Untargeted Moments: Rationalizing Facts 2 In our internal calibration, we focused on matching micro-level moments that reflect the lumpy nature of innovation, such as the average duration between spikes and the mean-max ratio of R&D expenses. For untargeted

moments, we concentrated on examining whether the model reproduces the aggregate characteristics of lumpy innovation. The last two rows of Table I 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 $s_{i,t}$.

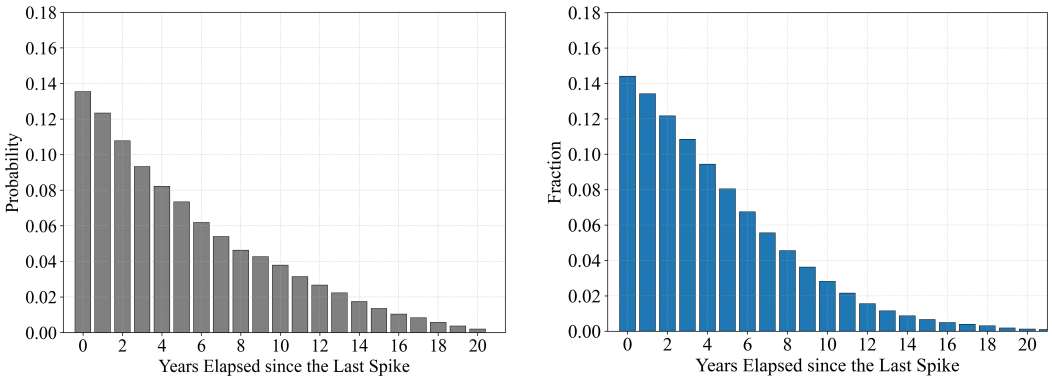
The following proposition overcomes this challenge by showing that there exists a one-to-one mapping from the distribution of unobserved ideas to the distribution of an observable: the time elapsed since the last spike. Intuitively, Proposition 2 holds because, in our model, an innovator who has waited a long time since the last spike adjustment in R&D expenses is more likely to have a big $s_{i,t}$.[8]

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 relationship holds:*

$$\mathcal{G}_t(s) = \sum_{n \in \mathcal{N}} P(s|TE_{i,t} = n)\mathcal{T}_t(n). \tag{29}$$

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.

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



(a) Data (Untargeted) (b) Model

FIGURE 5.—Distribution (Untargeted Moments): Time Elapsed since the Last Spike. Note: 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.

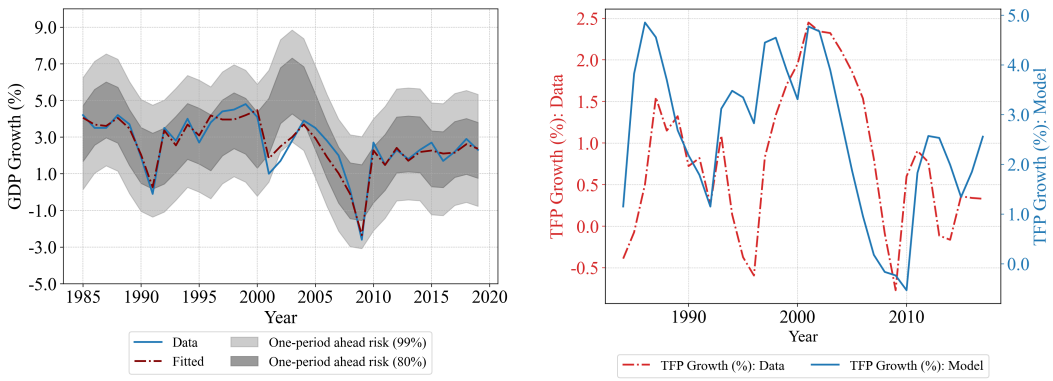
4.3. *Fitting the Model to the Data over Time*

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^{ϱ}.[9]

Figure 6a shows the fitting results, where the model-derived GDP growth rate closely matches the actual rate. Additionally, Figure 6a includes 80% and 95% probability bands for one-period ahead predictions, indicating that the GDP growth rate for period $t+1$ is likely to fall within this range based on the state at period t. 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 6b 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

[9]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.



(a) Model Fit (in-sample): GDP (b) Model Fit (out-of-sample): TFP

FIGURE 6.—Fitting Model to the Data over Time. Note: The TFP data is from the Federal Reserve Bank of San Francisco, which reports updated capital utilization-adjusted TFP using Fernald (2014)'s method. Because the TFP series contains other factors beyond technology growth (e.g., capital misallocation), the red line plots a three-period moving average of TFP.

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.

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 5.2 presents the main result of the paper, the Great Twins hypothesis. Section 5.3 introduces an additional prediction of the model: the recession trap. Model results in Section 5.2 and 5.3 are based on simulations.

5.1. *Novel Mechanism: Vulnerable Moderation*

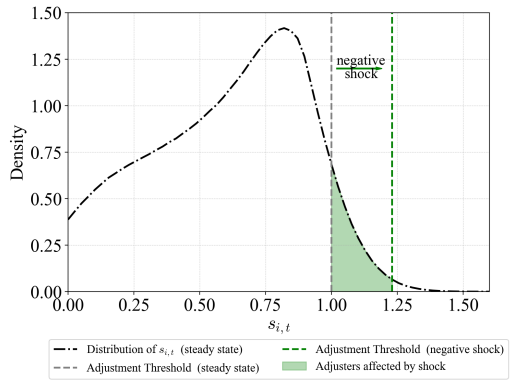
5.1.1. *The Asymmetric Effects of Shocks on the Distribution*

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.

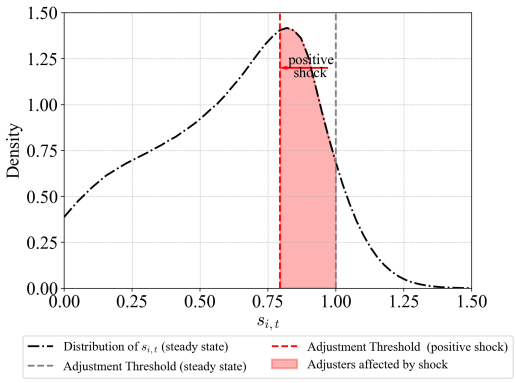
To see this, first consider a positive demand shock. The shock generates an economic boom, increasing the profitability of developing new varieties. This motivates innovators to more eagerly initiate new R&D projects, resulting in the leftward shift of the adjustment threshold, as indicated in Figure 7b 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 7c.

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 7a). 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 7c.

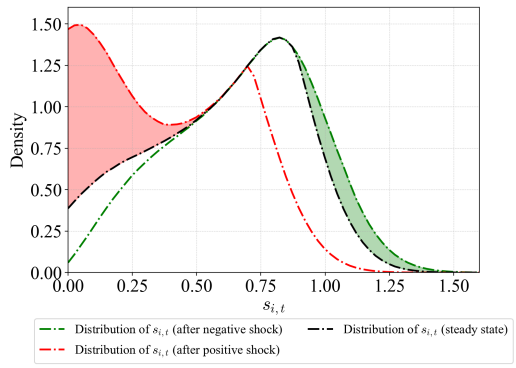
The stretched distribution due to a negative shock and the compressed distribution due to a positive shock are shown in Figure 7c. The following proposition shows that, under very general conditions, the compression from positive shocks (red shaded area) exceeds the stretching from negative shocks (green shaded area), given the presence of a fixed R&D cost.



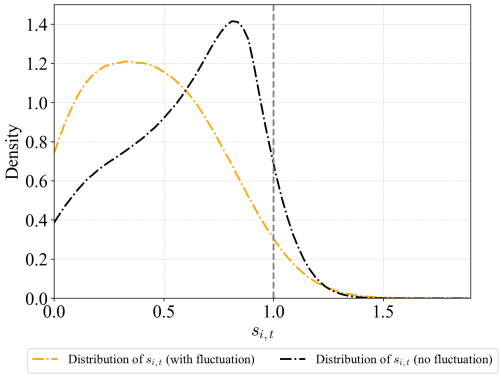
(a) Negative shocks shift $\bar{s}_{i,t}$ to the left



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



(c) Asymmetric Effects on the Distribution



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

FIGURE 7.—Illustration of Shocks and Effects on the Distribution of $s_{i,t}$

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 3.2. *Q.E.D.*

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, 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 $\overline{s}$ 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.

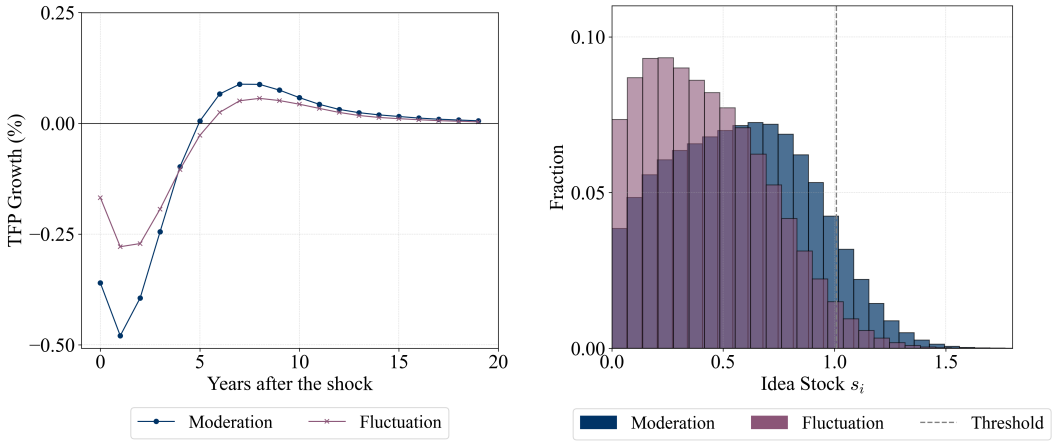
COROLLARY 1: *The threshold $\overline{c}$ described in Proposition 3 equals zero if the density function of the idea inflow, $\psi(\eta)$, decreases with η.*

PROOF: See Online Appendix 3.2. *Q.E.D.*

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 (or, equivalently, more likely to draw a mediocre idea than a great idea).

To further illustrate the critical role of lumpiness in deriving the model mechanism, in Online Appendix 2.3, we define asymmetry as the difference between the red and green areas in Figure 7c. We analytically show that, for two specific idea inflow processes, the degree of asymmetry increases with higher fixed adjustment costs (lumpiness). Our numerical analysis confirms this pattern across various idea inflow processes and model parameters.

Discussion: Relationship with Lumpy Capital Investment Models In Online Appendix 2.4, we further discuss 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.



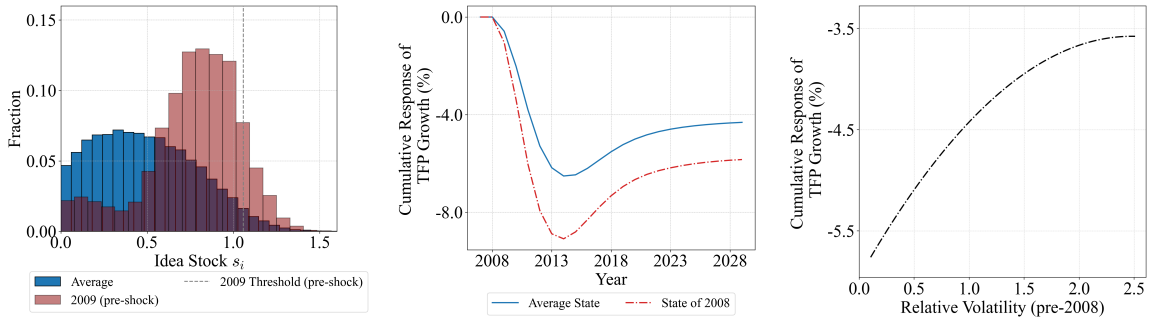
(a) State-Dependent Effects of Shocks (b) Mechanism: Distribution & Threshold

FIGURE 8.—The Impulse Response of Shocks in the Calibrated Model. Note: Figure 8a and 8b are based on simulations. The moderation (fluctuation) subsample is defined as panels that have the lowest (highest) 5% standard deviation of output growth in $[t_0 - 10, t_0)$. Figure 8b depicts the average distribution of $s_{i,t}$ in period t_0. Figure 8a plots the average impulse response of a one-standard-deviation negative demand shock at t_0.

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 7d. 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 8 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 8b. 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 8a.



(a) Distribution before the Great Recession

(b) CIR of TFP to the Great Recession Shocks

(c) Effects with Varying Moderation

FIGURE 9.—The Effect of Great Moderation. Note: Panel (a)-(c) are based on shocks and distributions of $s_{i,t}$ by fitting the model to actual GDP growth. The Great Recession shocks are defined as the 2008 and 2009 shocks. The average state is defined by having the same ϱ_t, though the distributions of $s_{i,t}$ may differ. In Figure 9c, ϱ_t in 1979-2002 corresponds to the fitted values from Section 4.3, while it simulates counterfactual ϱ_t for 2003-2007. Relative volatility is defined as the ratio of the 2003-2007 realized volatility of ϱ_t to its unconditional counterpart.

5.2. *Model Results: the Great Twins*

The results presented in Figure 8 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 6a).

Figure 9a 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 9b illustrates the decline in A_t resulting from the 2008 and 2009 recessionary shocks. The effects of the Great Recession shocks on A_t unfolded gradually, 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.

The results of the Cumulative Impulse Response (CIR) of A_t in the limit (from the year 2030 onward) indicate that the U.S. economy, having experienced the Great Recession, now exhibits a GDP nearly 6% lower than a counterfactual U.S. economy without the

Great Recession shocks; see the red dashed line.[10] 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 4% in the limit; see the blue solid line. This implies that about 1/3 of the total economic loss can be attributed to the distribution of idea stock before the Great Recession.

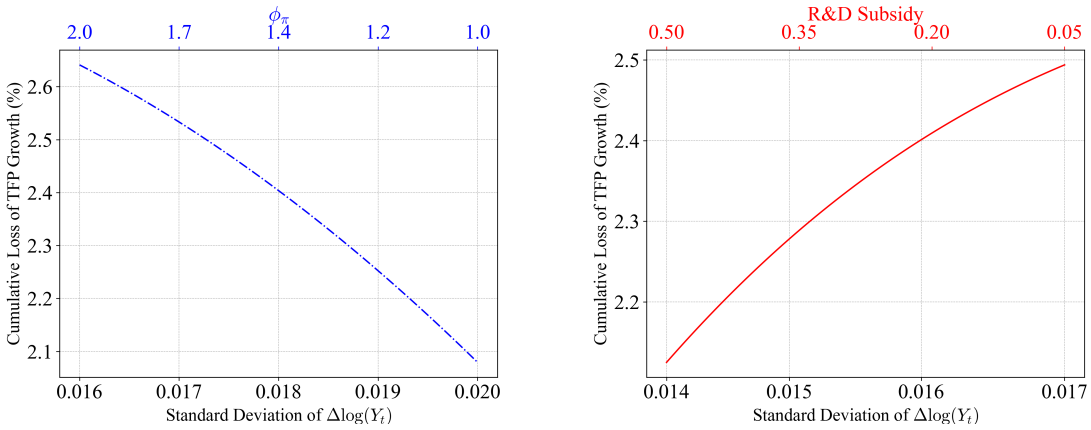
Next, we pose the question: What contributed to the 2009 pre-shock distribution of the idea stock, which displays a significant mass around the extensive margin? We assert that this pre-shock distribution, implying a notable left-tail risk, was caused by the moderated economic fluctuations during the 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 9c 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 percentage points.

The above analysis together implies that about 1/3 of the total economic loss, measured by the permanent loss in real GDP in the limit, can be attributed to the additional vulnerabilities created by the Great Moderation. Alternatively, the Great Recession is 50% more severe due to the preceding moderation——the main emphasis of the paper.

It is worth clarifying that reduced sample volatility does not imply reduced risk. Specifically, the changes along the x-axis in Figure 9c 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.

Discussion: Ideas are Getting Harder to Find Recent empirical evidence (Bloom et al. 2020, Kogan et al. 2017) suggests that ideas have been getting harder to find over the past 40 years. How does this characteristic affect our results? In Online Appendix, we show that for a lower average inflow of ideas, the 2008 Great Recession caused a greater loss of growth. This suggests an increasing importance of the mechanism/insight emphasized in our paper moving forward.

[10]Note that this does not imply that the GDP is 6% lower than its pre-shocks level since the counterfactual U.S. economy without the Great Recession shocks would have experienced positive growth.



(a) The Tail-Risk Variance Trade-off

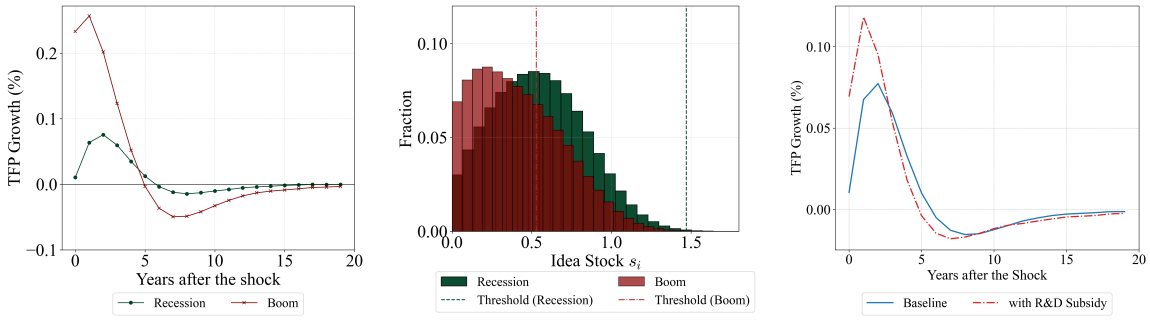
(b) Splitting the Great Twins

FIGURE 10.—The Great Twins: Policy Implications. 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 cumulative loss in TFP is measured as the cumulative impulse response over 40 periods following the shock.

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 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.



(a) State-Dependent Effects of Shocks (b) Mechanism: Distribution & Threshold (c) Escaping the Recession Trap

FIGURE 11.—The Recession Trap and Escape Mechanisms. Note: Panel (a) - (c) are based on simulations. At t_0, a boom is defined as $\varrho_{t_0} = -\text{std}(\varrho_t)$, and a recession is defined as $\varrho_{t_0} = \text{std}(\varrho_t)$. Panel (a) depicts the average distribution of s_{i,t_0}. Panel (b) plots the average impulse response to a monetary shock of 25-basis points in period t_0 of the boom and recession subsamples. Panel (c) plots the average impulse response in the recession state with and without an R&D subsidy amounting to 40% of the fixed adjustment cost.

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 mitigating 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 11a 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 $\varepsilon_{t+1}^m = \rho_m \varepsilon_t^m + e_{t+1}^m$.

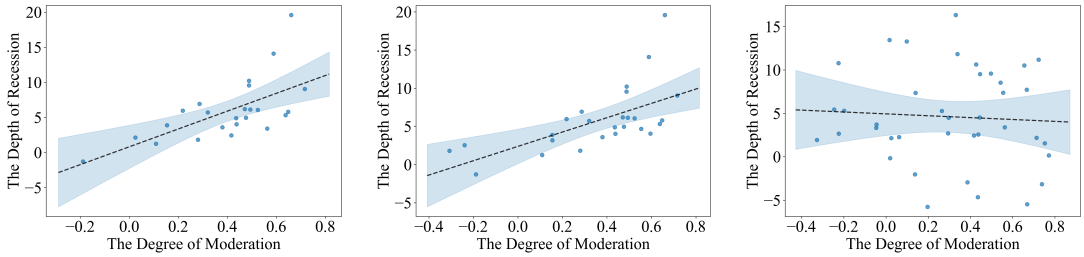
Figure 11a shows that the effects of monetary policy during recessions, periods where the effects of stimulus policy are most appreciated, are substantially dampened compared to the effects in a boom state. At their peak, the effects in a recession state are reduced to one-third of the effects in a boom state.

Figure 11b 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 11c 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.

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. Section 6.3 demonstrates that the Great Recession increases the average quality of innovations, a natural outcome of the extensive margin mechanism.



(a) OECD Countries (b) R&D-to-GDP > 1% (c) R&D-to-GDP < 1%

FIGURE 12.—The Great Twins in the Data. Notes: The result is based on the cross-country data from the World Bank. To avoid the depth of recession being influenced by the pre-recession trend, we normalize the depth of recession by the pre-recession trend growth, defined as the average growth rate over the past fifteen years. The correlation in Figure 12a is 0.64 ($p = 0.00$), the correlation for the countries with R&D-to-GDP ratio greater than 1% is 0.61 ($p = 0.00$), and for the countries with R&D-to-GDP ratio smaller than 1% is -0.06 ($p = 0.70$)

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 moderation? 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.[11]

Figure 12a shows that the depth of recessions positively correlates with the degree of moderation in the pre-crisis periods for OECD countries, consistent with the Great Twins hypothesis. Figure 12b demonstrates this relationship for countries with an R&D-to-GDP ratio greater than one percent[12].

[11] See Online Appendix 1.4 for details of the formal definition of the depth of recession.

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

Our theory describes the medium-term business cycle fluctuations of advanced economies that conduct innovation; it does not fit developing economies or economies that do not conduct R&D. Consistent with this prediction, we find that for countries with an R&D-to-GDP ratio below 1%, the correlation between the degree of moderation and recession depth disappears: -0.06 ($p = 0.70$). See Online Appendix 1.4 for a more detailed robustness check of the sample selection rule.

Robustness Checks In Online Appendix 1.4, we demonstrate that the Great Twins relationship persists even when controlling for financial factors. Additionally, we find that the model-implied relationship between recessions and their preceding moderation is evident when considering TFP growth (instead of real GDP) and across other recessions (beyond the 2008 recession), using data from 18 advanced economies since World War II. We conclude that the Great Twins is a robust phenomenon across countries and over time.

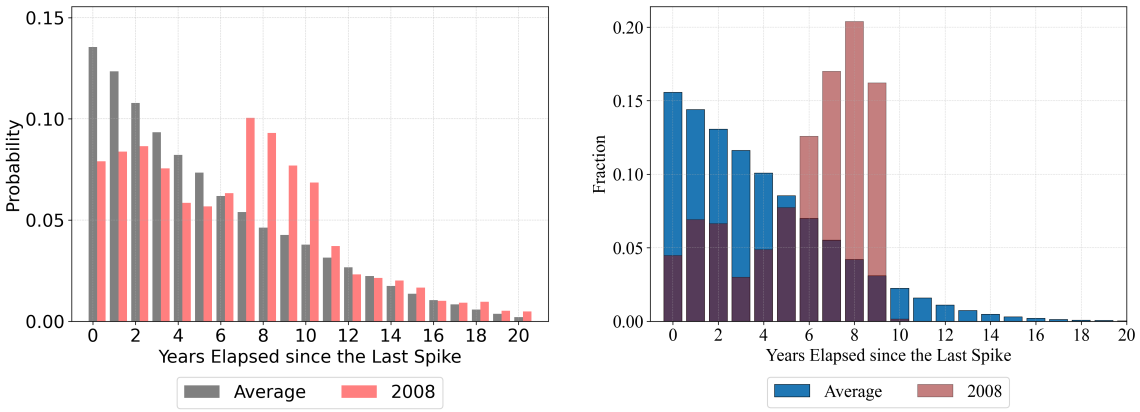
6.2. *Micro-Evidence: Distribution in the Data*

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 9a, 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 13 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 13b, 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 13a, this pattern is also mirrored in the empirical data.

6.3. *Patent Quality over the 2008 Recession*

One implication of the model is that recessions increase the adjustment threshold, meaning that only research projects containing better ideas will be set up. Therefore, after a re-



(a) Distribution: in the Data (b) Distribution: in the Model

FIGURE 13.—Distribution of Time Elapsed since the Last Spike: Pre-2008 Recession. Note: Figure 13a 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 13b 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.

cession, we should observe an increase in the *average* patent quality. In Figure 14, we plot the time series of the average patent quality after 1997.[13] We observe that patent quality shows a declining trend before 2008, with the trend *reversing after the recession*, suggesting an increase in patent quality due to the Great Recession.

7. CONCLUSION

We document three novel facts. First, R&D investment decisions are lumpy: innovating 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 in-

[13] In R&D literature, patent citations are a common measure of patent quality. Since our data only goes up to 2022, citations for patents post-2008 are not stationary. Thus, we define average patent quality as the probability of patents being cited within five years of issuance.

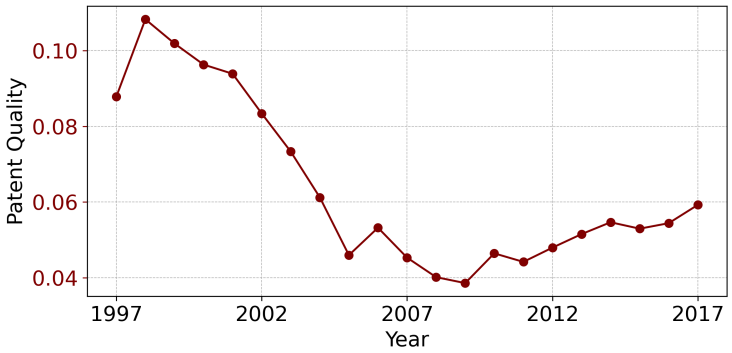


FIGURE 14.—Patent Quality Overtime. Notes: The data used for plotting is sourced from the United States Patent and Trademark Office (USPTO). Patent quality is defined as the rate at which patents issued in a given year are cited by newly issued patents within the following five years.

novation 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, uncover micro-level evidence supporting the underlying mechanism, and document the effects of the Great Recession on the average quality of patents, consistent with the highlighted extensive margin mechanism.

Overall, we view the main contribution of our paper as providing a new qualitative insight: moderation can causally contribute to a more severe recession. Quantitatively, according to our calibration, the Great Recession was 50% more severe due to the preceding Great Moderation. However, if one is interested in the sources of the Great Recession, a more thorough quantitative analysis that compares our mechanism with existing ones can be a fruitful path to take in future research.

REFERENCES

ACEMOGLU, DARON, PHILIPPE AGHION, LEONARDO BURSZTYN, AND DAVID HEMOUS (2012): "The environment and directed technical change," *American economic review*, 102 (1), 131–166. [19]

ACEMOGLU, DARON, UFUK AKCIGIT, HARUN ALP, NICHOLAS BLOOM, AND WILLIAM KERR (2018): "Innovation, reallocation, and growth," *American Economic Review*, 108 (11), 3450–91. [6]

AGHION, PHILIPPE AND PETER HOWITT (1992): "A Model of Growth Through Creative Destruction," *Econometrica*. [6]

AKCIGIT, UFUK AND WILLIAM R KERR (2018): "Growth through heterogeneous innovations," *Journal of Political Economy*, 126 (4), 1374–1443. [6]

ANZOATEGUI, DIEGO, DIEGO COMIN, MARK GERTLER, AND JOSEBA MARTINEZ (2019): "Endogenous technology adoption and R&D as sources of business cycle persistence," *American Economic Journal: Macroeconomics*, 11 (3), 67–110. [7, 25, 29]

BACHMANN, RÜDIGER, RICARDO J CABALLERO, AND EDUARDO M R A ENGEL (2013): "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model," *American Economic Journal: Macroeconomics*, 5 (4), 29–67. [6]

BALEY, ISAAC AND ANDRÉS BLANCO (2021): "Aggregate dynamics in lumpy economies," *Econometrica*, 89 (3), 1235–1264. [6]

BALL, LAURENCE (2014): "Long-term damage from the Great Recession in OECD countries," *European Journal of Economics and Economic Policies: Intervention*, 11 (2), 149–160. [39]

BARLEVY, GADI (2007): "On the cyclicality of research and development," *American Economic Review*, 97 (4), 1131–1164. [6, 11]

BENIGNO, GIANLUCA AND LUCA FORNARO (2018): "Stagnation traps," *The Review of Economic Studies*, 85 (3), 1425–1470. [7]

BIANCHI, FRANCESCO, HOWARD KUNG, AND GONZALO MORALES (2019): "Growth, slowdowns, and recoveries," *Journal of Monetary Economics*, 101, 47–63. [7]

BLOOM, NICK (2007): "Uncertainty and the Dynamics of R&D," *American Economic Review*, 97 (2), 250–255. [6, 18]

BLOOM, NICHOLAS, CHARLES I. JONES, JOHN VAN REENEN, AND MICHAEL WEBB (2020): "Are Ideas Getting Harder to Find?" *American Economic Review*, 110 (4), 1104–1144. [29, 35]

BRODA, CHRISTIAN AND DAVID E WEINSTEIN (2010): "Product creation and destruction: Evidence and price implications," *American Economic Review*, 100 (3), 691–723. [6]

BROWN, JAMES R AND BRUCE C PETERSEN (2011): "Cash holdings and R&D smoothing," *Journal of Corporate Finance*, 17 (3), 694–709. [6, 18]

CABALLERO, RICARDO J AND EDUARDO MRA ENGEL (1999): "Explaining investment dynamics in US manufacturing: a generalized (S, s) approach," *Econometrica*, 67 (4), 783–826. [6]

CABALLERO, RICARDO J, EDUARDO MRA ENGEL, JOHN C HALTIWANGER, MICHAEL WOODFORD, AND ROBERT E HALL (1995): "Plant-level adjustment and aggregate investment dynamics," *Brookings papers on economic activity*, 1–54. [6]

CETTE, GILBERT, JOHN FERNALD, AND BENOIT MOJON (2016): "The pre-Great Recession slowdown in productivity," *European Economic Review*, 88, 3–20. [29]

CHEN, ZHAO, ZHIKUO LIU, JUAN CARLOS SUÁREZ SERRATO, AND DANIEL YI XU (2021): "Notching R&D investment with corporate income tax cuts in China," *American Economic Review*, 111 (7), 2065–2100. [6, 18]

CHIAVARI, ANDREA AND SAMPREET SINGH GORAYA (2024): "The Rise of Intangible Capital and the Macroeconomic Implications," *Working Paper*. [6]

COMIN, DIEGO AND MARK GERTLER (2006): "Medium-term business cycles," *American Economic Review*, 96 (3), 523–551. [2, 3, 5, 7, 12, 25]

COMIN, DIEGO A, MARK GERTLER, AND ANA MARIA SANTACREU (2009): "Technology innovation and diffusion as sources of output and asset price fluctuations," Tech. rep., National Bureau of Economic Research. [5, 12]

COOPER, RUSSELL W AND JOHN C HALTIWANGER (2006): "On the nature of capital adjustment costs," *The Review of Economic Studies*, 73 (3), 611–633. [11]

CÚRDIA, VASCO, ANDREA FERRERO, GING CEE NG, AND ANDREA TAMBALOTTI (2015): "Has US monetary policy tracked the efficient interest rate?" *Journal of Monetary Economics*, 70, 72–83. [22]

FABRIZIO, KIRA R AND ULYA TSOLMON (2014): "An empirical examination of the procyclicality of R&D investment and innovation," *Review of Economics and Statistics*, 96 (4), 662–675. [6]

FATAS, ANTONIO (2000): "Do business cycles cast long shadows? Short-run persistence and economic growth," *Journal of economic growth*, 5, 147–162. [6]

FERNALD, JOHN (2014): "A quarterly, utilization-adjusted series on total factor productivity," Federal Reserve Bank of San Francisco. [28, 29]

GALÍ, JORDI AND LUCA GAMBETTI (2009): "On the sources of the great moderation," *American Economic Journal: Macroeconomics*, 1 (1), 26–57. [36]

GNEWUCH, MATTHIAS AND DONGHAI ZHANG (2023): "Monetary policy, firm heterogeneity, and the distribution of investment rates," *SSRN Working Paper*. [6]

GORNEMANN, NILS, PABLO GUERRÓN-QUINTANA, AND FELIPE SAFFIE (2021): "Real Exchange Rates and Endogenous Productivity," . [5]

GRILICHES, ZIV (1990): "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, 28 (4), 1661–1707. [7]

GUERRON-QUINTANA, PABLO A AND RYO JINNAI (2019): "Financial frictions, trends, and the great recession," *Quantitative Economics*, 10 (2), 735–773. [7]

HSIEH, CHANG-TAI AND PETER J KLENOW (2009): "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 124 (4), 1403–1448. [29]

JUSTINIANO, ALEJANDRO AND GIORGIO E PRIMICERI (2008): "The time-varying volatility of macroeconomic fluctuations," *American Economic Review*, 98 (3), 604–641. [36]

KHAN, AUBHIK AND JULIA THOMAS (2008): "IDIOSYNCRATIC SHOCKS AND THE ROLE OF NONCONVEXITIES," *Econometrica*, 76 (2), 395–436. [6]

KLETTE, TOR JAKOB AND SAMUEL KORTUM (2004): "Innovating firms and aggregate innovation," *Journal of political economy*, 112 (5), 986–1018. [6]

KOBY, YANN AND CHRISTIAN K WOLF (2020): "Aggregation in Heterogeneous-Firm Models: A Sufficient Statistics Approach *," Tech. rep. [6]

KOGAN, LEONID, DIMITRIS PAPANIKOLAOU, AMIT SERU, AND NOAH STOFFMAN (2017): "Technological innovation, resource allocation, and growth," *The quarterly journal of economics*, 132 (2), 665–712. [35]

KRUSELL, PER AND ANTHONY A SMITH, JR (1998): "Income and wealth heterogeneity in the macroeconomy," *Journal of political Economy*, 106 (5), 867–896. [23]

KUNG, HOWARD AND LUKAS SCHMID (2015): "Innovation, growth, and asset prices," *The Journal of Finance*, 70 (3), 1001–1037. [5, 25]

LIPPI, FRANCESCO AND ALEKSEI OSKOLKOV (2023): "A Structural Model of Asymmetric Lumpy Investment," . [6]

MATSUYAMA, KIMINORI (1999): "Growing through cycles," *Econometrica*, 67 (2), 335–347. [5]

MIAO, JIANJUN AND PENGFEI WANG (2014): "A Q-theory model with lumpy investment," *Economic theory*, 57, 133–159. [6]

MOSCARINI, GIUSEPPE AND FABIEN POSTEL-VINAY (2018): "The cyclical job ladder," *Annual Review of Economics*, 10, 165–188. [29]

OTTONELLO, PABLO AND THOMAS WINBERRY (2020): "Financial heterogeneity and the investment channel of monetary policy," *Econometrica*, 88 (6), 2473–2502. [24]

OUYANG, MIN (2011): "On the Cyclicality of R&D," *Review of Economics and Statistics*, 93 (2), 542–553. [6]

QUERALTO, ALBERT (2020): "A model of slow recoveries from financial crises," *Journal of Monetary Economics*, 114, 1–25. [7]

ROMER, PAUL M (1990): "Endogenous technological change," *Journal of political Economy*, 98 (5, Part 2), S71–S102. [3, 5, 14, 15]

SCHMITZ, TOM (2021): "Endogenous growth, firm heterogeneity and the long-run impact of financial crises," *European Economic Review*, 132, 103637. [7]

WICKSELL, KNUT (1898): "Interest and Prices (English translation by RF Kahn, 1936)," . [23]

WINBERRY, THOMAS (2021): "Lumpy investment, business cycles, and stimulus policy," *American Economic Review*, 111 (1), 364–96. [6]

ZHANG, DONGHAI (2022): "Term structure, forecast revision, and the signaling channel of monetary policy," *Journal of the European Economic Association*, 20 (4), 1522–1553. [22]