Lumpy to Slumpy: How Financing Costs drive Slow Recoveries *

Miguel H. Ferreira † Timo Haber ‡ Hanbaek Lee §

June 16, 2026

Abstract

Why do some recessions lead to slow recoveries? We argue that recovery speed depends on how financial fragility interacts with lumpy investment. Large adjustments require external finance and weaken balance sheets, which in turn make future adjustment less attractive. This feedback, the *lumpy-investment spiral*, lengthens non-adjustment spells and raises dispersion in the time since firms last adjusted. In standard lumpy-adjustment models, older non-adjusters tend to close capital gaps quickly because they are near adjustment triggers. The spiral reverses this force: delayed adjustment depletes internal resources, raises external-finance needs, and leaves large gaps among slow-adjusting firms. In a calibrated heterogeneous-firm business-cycle model, the spiral force dominates and substantially slows recoveries. Firm-level evidence supports both sides of the mechanism and shows that weak recoveries are preceded by many financially fragile firms deep into non-adjustment spells.

Keywords : Investment, Financial Frictions, Firm Dynamics
JEL Codes: D53, E44, D21

*We thank our discussants Marcel Peruffo and John Finlay. For insightful comments and discussions we thank, Isaac Baley, Vasco M. Carvalho, Julieta Caunedo, Andres Erosa, Miguel Faria-e-Castro, Priit Jeenas, Julian Kozlowski, Tim Lee, Francesco Lippi, Ramon Marimon, Davide Melcangi, Pascal Paul, Alessandra Peter, Ricardo Reis, Edouard Schaal, Tom Schmitz, Tatsuro Senga, Sanjay R. Singh, Vincent Sterk, Christian Wolf. For questions and comments, we also thank seminar participants at U.C.Davis, U.Kent U.Manchester, QMUL SEF, the BSE summer forum "Macro Fluctuations with Micro Frictions" workshop, the SED (Cartagena), the EEA (Barcelona), Lisbon Macro Workshop, LuBraMacro (Porto), the PEJ (Braga), the MMF (Portsmouth), CICM (Hong Kong), International Conference on European Studies (Zurich), Midwest Macro (Kansas City), T2M (Paris). Previously circulated as "From Spreads to Spirals: How Financial Frictions Drive Lumpy Investments"

†Queen Mary University of London. migueltchferreira@gmail.com

‡De Nederlandsche Bank, Eurosystem. t.f.haber@dnb.nl. Views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank or the Eurosystem.

§University of Cambridge. hanbaeklee1@gmail.com

1 Introduction

Why are some recessions followed by slow recoveries, as in the aftermath of the 2008 financial crisis? While the nature of the initial shock clearly matters, recessions with similar aggregate declines can be followed by very different recovery paths. We argue that these dynamics are shaped by the joint distribution of firms' financial capacity and accumulated investment needs. Slow recoveries arise when many firms enter a downturn with large capital gaps but limited internal resources, so that closing these gaps requires costly external finance. We show that embedding this interaction between financial fragility and lumpy investment in a business-cycle model generates heterogeneous recovery speeds following recessions.

This paper develops this argument in three steps. First, we propose a micro-level mechanism through which fixed adjustment costs and financial frictions jointly generate prolonged investment non-adjustment spells. Large investment spikes require external finance and weaken firms' balance sheets; weaker balance sheets then reduce the probability of future adjustment, generating longer non-adjustment spells and more concentrated investment. We label this feedback loop the lumpy-investment spiral.

Second, we show that the lumpy-investment spiral generates state-dependent recovery dynamics by reshaping the distribution of firms over non-adjustment spells. The spiral lengthens these spells and increases dispersion in the time since firms last adjusted, leaving more firms with large accumulated capital gaps. In a standard lumpy-adjustment model, this distributional shift would tend to speed recovery: firms that have gone longest without adjusting are closer to their adjustment triggers and are therefore more likely to close their gaps. With financial frictions, this stabilizing age gradient is weakened or reversed. Firms with longer spells have depleted internal resources, so closing larger capital gaps requires more external finance and becomes more costly. Quantitatively, this financial channel dominates the standard trigger effect. When a downturn hits an economy with many firms in this slow-moving tail, aggregate capital gaps close more slowly and recoveries become more persistent.

Third, we provide empirical support for both sides of the spiral mechanism and for its aggregate implications. In firm-level data, tax-induced investment spikes worsen firms' financial positions, while firms with weaker balance sheets are less likely to spike. At the aggregate level, recessions followed by weak recoveries are preceded by a larger mass of firms with both weak financial positions and long non-adjustment spells, as well as greater dispersion in non-adjustment durations. These are precisely the distributional conditions under which the model predicts slow capital recovery.

We begin by showing why these conditions are empirically relevant. Using annual Compustat data on U.S. publicly listed firms, we document three robust patterns linking firms' financial positions, lumpy investment behavior, and subsequent recovery speeds. First, recessions that are followed by slow recoveries are preceded by weaker firm financial positions and longer investment non-adjustment spells.[1] Second, investment lumpiness is positively associated with tighter financial constraints in the cross section. Firms with lumpier investment profiles tend to have lower distance to default and greater reliance on external finance. Third, the extensive margins of debt issuance and investment move together over time: the fraction of firms with debt spikes consistently tracks the fraction with investment spikes. These facts suggest a close link between investment and financing decisions and the speed of recoveries following recessions, and motivate the central hypothesis of this paper: large investment events are associated with higher financing costs.

Guided by this hypothesis, we develop a two-period model that incorporates both fixed adjustment costs and costly external finance. The model clarifies how the financing cost of large investment events can itself create non-adjustment regions and amplify real adjustment frictions. In this framework, investment spikes occur when firms cross a non-adjustment threshold, incurring both real and financial adjustment costs. Financial frictions alone can generate non-adjustment spells, even in the absence of real adjustment costs. This occurs when there is a discontinuity between the internal and external cost of funds: if the marginal benefit of an additional unit of investment lies between the marginal cost of internal and external financing, a firm with insufficient internal resources to expand its capital stock remains in a non-adjustment region.

Crucially, the interaction between real and financial frictions gives rise to a self-reinforcing mechanism that we term the lumpy-investment spiral. A firm that relies heavily on external financing may find it unprofitable to pay the adjustment cost and thus postpone investment. As investment is delayed, internal resources depreciate, increasing the cost of external finance needed to reach the same capital level. The higher financing cost then feeds back into both the extensive and intensive investment margins, completing the feedback loop.

We then ask whether this firm-level feedback is strong enough to shape aggregate recovery dynamics. To do so, we embed this mechanism in a business-cycle general equilibrium model with heterogeneous firms that face idiosyncratic productivity shocks, capital adjustment frictions, and costly external finance. Firms choose capital and external

[1] We follow the definition of Donaldson & Wieland (2025) and label slow recoveries when post-recession growth is weaker than pre-recession growth.

finance optimally, facing fixed costs of adjusting capital and a spread on external borrowing. The model is solved numerically and calibrated to match key U.S. firm-level moments, including the frequency and size of investment spikes, the average external financing share, and labor input.

The calibrated model shows that the spiral changes both the allocation of capital across firms and the speed at which aggregate capital recovers after a downturn. To connect these micro-level distortions to aggregate recovery speed, we follow Baley & Blanco (2021) and summarize the relevant cross-section of firms. The key object is how a firm's capital-to-productivity gap, $x \equiv \log(k/z)$, lines up with the time a since its last adjustment: when larger gaps are concentrated among longer-inactive firms, so that the covariance $\mathrm{Cov}(x, a)$ is more negative, aggregate capital reverts more slowly. The lumpy-investment spiral changes precisely this cross-sectional object in two ways. First, the spiral misallocates capital. The cross-sectional correlation between productivity and capital falls from 0.84 in a frictionless economy to 0.61 in the baseline, below the 0.67 obtained in an otherwise identical economy without financial frictions. Thus, costly external finance worsens capital allocation precisely when firms also face lumpy adjustment costs: productive firms remain stuck below their desired scale, while financial frictions have little aggregate effect in isolation.

Second, and central to the slow recovery, the spiral shifts firms toward longer non-adjustment durations and attenuates the standard self-correcting force that the age distribution exerts on aggregate convergence. In a standard lumpy economy, firms that have gone longest without adjusting are closest to their adjustment triggers, so the age distribution is stabilizing: aggregate capital gaps close *faster* because the firms with the largest gaps are also the most likely to adjust. Costly external finance instead lowers the adjustment probability for firms deep into their spells: long non-adjustment spells leave firms with lower internal resources and higher external-finance needs. This broadens the dispersion of non-adjustment durations. The average time since a firm last adjusted rises from 4.5 to 5.3 years, and its variance rises from 16.5 to 21.5, so capital shortfalls become concentrated among the slowest-adjusting firms. As a result, the covariance between the capital-to-productivity gap and non-adjustment duration falls to -0.99, compared with -0.87 without financial frictions. Because the cumulative response of aggregate capital is governed by this covariance, costly external finance lengthens the capital half-life by roughly 14% relative to lumpy adjustment alone, while financial frictions in isolation have no impact on the capital half-life.

The source of this aggregate persistence is highly uneven across firms. The nonlinear aggregate effects stem primarily from high-productivity firms, which have the highest marginal benefit of investment and rely most on external finance. Their marginal benefit curve is also the most responsive to negative TFP shocks, an effect that is further amplified when internal resources are scarce. In low-internal-resource states, a larger mass of firms is also close to non-adjustment and spiral regions. A negative productivity shock then compresses profits and pushes more firms into the spiral, causing recessions to be more persistent. Shutting down either financial or real frictions eliminates this state dependence, as micro-level nonlinearities are no longer strong enough to survive aggregation.

Finally, we return to the data and provide evidence of the four links that connect firms' investment and financial decisions to aggregate recovery speed. First, large investment adjustments coincide with the use of external finance and drawdown of internal liquidity. Second, we exploit variation in investment spikes induced by firms' exposure to positive bonus-depreciation changes and find that tax-induced investment spikes generate a large and statistically significant deterioration in distance to default on impact. This result directly supports the first direction of the mechanism: undertaking a large investment adjustment weakens firms' balance sheets. Third, we show that financially weaker firms are less likely to undertake subsequent investment spikes. In a linear probability model with firm and year fixed effects, the estimated spike coefficients decline as firms become more constrained. This evidence supports the feedback direction of the mechanism: lumpy investment weakens balance sheets, which delays future lumpy adjustment. Finally, we connect the firm-level mechanism to aggregate recovery dynamics. Recessions followed by weak recoveries are preceded by substantially weaker firm balance sheets and by a larger mass of firms jointly characterized by low distance to default and long investment non-adjustment spells. In addition, the variance of the non-adjustment spells is higher before weak recoveries. These two channels combined are the exact mechanism identified by the model that leads to slower recoveries.

Taken together, the model and evidence point to a distributional source of slow recoveries. Recovery dynamics are not governed only by the size of the initial shock, average financial fragility, or average capital gaps, but by whether firms with accumulated investment needs also lack the internal resources to adjust. This perspective helps explain why recessions with similar aggregate declines can lead to different recovery paths: downturns are more persistent when they hit an economy in which the firms that most need to invest are also the least able to finance adjustment.

Literature. Our paper contributes first to the literature on lumpy investment and aggregate dynamics. A large body of work shows that plant- and firm-level investment is infrequent and concentrated in large adjustment episodes (Cooper et al., 1999; Cooper &

Haltiwanger, 2006), and studies how non-convex adjustment costs shape aggregate investment dynamics (Bachmann et al., 2013; Caballero & Engel, 1999; Gnewuch & Zhang, 2025; Gourio & Kashyap, 2007; Khan & Thomas, 2008; Lee, 2022; Thomas, 2002; Winberry, 2021). Most closely, Baley & Blanco (2021, 2024) show that the aggregate consequences of lumpy adjustment can be summarized by cross-sectional sufficient statistics, including how capital gaps covary with the time since last adjustment. We build directly on this insight. Our contribution is to show that financial fragility changes these sufficient-statistic objects themselves. The lumpy-investment spiral lengthens non-adjustment spells, increases dispersion in adjustment ages, and places large capital gaps among firms that adjust slowly. Thus, finance matters not simply by adding another wedge to a lumpy economy, but by reshaping the cross-sectional state that governs aggregate convergence.

Second, the paper contributes to the literature on financial frictions and firm investment. Canonical financial-accelerator and heterogeneous-firm models show that balance-sheet conditions, borrowing costs, and liquidity constraints shape investment and aggregate fluctuations (Begenau & Salomao, 2018; Bernanke et al., 1999; Buera & Karmakar, 2022; Cloyne et al., 2023; Crouzet, 2018; Jeenas, 2023; Jermann & Quadrini, 2012; Khan & Thomas, 2013; Kiyotaki & Moore, 1997; Ottonello & Winberry, 2020; Xiao, 2022). Related work studies environments in which financing conditions interact with real adjustment margins, including Senga et al. (2017), Jiao & Zhang (2022), and Melcangi (2024). Relative to this literature, we isolate a specific dynamic interaction between real and financial frictions. Large lumpy adjustments are externally-finance intensive and can weaken firms' subsequent financial positions; weaker balance sheets then reduce the likelihood of future adjustment. This feedback makes financial frictions powerful precisely when firms also face lumpy real adjustment costs. In the calibrated model, financial frictions alone have small aggregate effects, but their interaction with lumpy adjustment generates misallocation, longer non-adjustment spells, and slower recoveries. Earlier empirical evidence by Whited (2006) links financial constraints to investment lumpiness. We add evidence on both directions of the spiral: large investment spikes deteriorate firms' financial positions, and weaker financial positions subsequently predict lower spike probabilities, conditional on the time since the last spike.

Third, the paper relates to work on why some recessions are followed by weak recoveries. Existing explanations emphasize bad luck (Galí et al., 2012), investment boom-bust cycles (Beaudry et al., 2018; Rognlie et al., 2018), credit boom-bust cycles (Jordà et al., 2013), differences in sectoral shock incidence (Beraja & Wolf, 2022), changes in beliefs (Kozlowski et al., 2019, 2020), secular structural change (Leamer, 2021), hysteresis (Benigno & Fornaro, 2018; Garga & Singh, 2021; Reifschneider et al., 2015), and exogenous trend slowdowns (Fernald et al., 2025). Donaldson & Wieland (2025) emphasize that distinguishing among these explanations is difficult in aggregate data because the relevant counterfactual output path is unobserved. We take a complementary approach. Rather than inferring an unobserved aggregate counterfactual, we take recovery classifications to the firm distribution and ask whether weak recoveries are preceded by the distributional state that the model predicts should slow capital recovery. The relevant state is not only average firm financial fragility, nor only the distance of capital from its desired level, but the joint distribution of financial fragility and accumulated investment needs. The empirical exercises show that this distributional state is present before weak U.S. recoveries, and the quantitative model maps it into slower aggregate capital recovery.

Outlook The rest of the paper is structured as follows. Section 2 presents the stylized facts. Section 3 describes the basic theory. This is followed by Section 4, which describes the model before presenting the quantitative results in Section 5. Section 6 presents empirical support for the spiral mechanism. Finally, Section 7 concludes.

2 Stylized Facts

In this section, we use annual firm-level data from Compustat and CRSP to document three stylized facts that motivate the model. The organizing idea is that weak recoveries depend not only on firms' investment gaps, but also on their ability to finance the large adjustments needed to close them. We show that weak-recovery recessions are preceded by weaker balance sheets and longer time since the last investment spike relative to other recessions. We then show that this real-financial connection is present more broadly: firms with more concentrated investment histories are financially more fragile, and years with many investment spikes are also years with many debt issuance spikes. These patterns suggest that large real adjustments are financially intensive and motivate a model in which non-convex investment adjustment interacts with financial frictions.

Data and measurement. We use annual firm-level data from Compustat merged with CRSP. Our main real variable is the investment rate, defined as capital expenditures, Compustat item CAPX, divided by lagged property, plant, and equipment, item PPEGT, following Ottonello & Winberry (2024). We define an investment spike as an investment rate above 20 percent, in line with the lumpy-investment literature (see e.g. Bachmann & Ma, 2016; Cooper & Haltiwanger, 2006). Our main measure of financial position is distance to default, constructed using CRSP stock price data and the methodology of Merton (1974)

Table 1: Firm conditions before weak and non-weak recoveries: percentage deviations

Recovery group	Years since last investment spike	Distance to default
Non-weak recoveries	-18.5%	+18.0%
Weak recoveries	-6.1%	-6.5%

Note: The entries report percentage deviations from the full firm-year sample average for each variable. The non-weak-recovery group includes the 1969–70, 1973–75, and 1980–82 recessions. The weak-recovery group includes the 1990–91, 2001, and 2007–09 recessions. The 2020 pandemic recession is classified as a non-weak recovery, but is excluded from this headline comparison because its shutdown-and-reopening dynamics make it a special episode relative to standard business-cycle recessions.

and Gilchrist & Zakrajšek (2012). Lower distance to default indicates greater financial fragility and a higher implied external finance premium, consistent with evidence that distance to default negatively predicts actual bond spreads (see e.g. Ferreira et al., 2024; Gilchrist & Zakrajšek, 2012). Appendix A.1 provides further details on data cleaning, variable construction, and descriptive statistics.

2.1 Fact 1: Weak recoveries follow delayed adjustment and fragile balance sheets

We begin by comparing pre-recession firm conditions across weak- and non-weak-recovery episodes. We follow the classification in Donaldson & Wieland (2025), which labels the 1990–91, 2001, and 2007–09 recessions as weak-recovery episodes, and the 1969–70, 1973–75, and 1980–82 recessions as non-weak-recovery episodes. For each recession, we measure firm conditions in the calendar year before the recession starts.[2,3]

Table 1 summarizes the pre-recession state of firms in the two recovery groups, with entries reported as percentage deviations from the corresponding averages. Firstly, on the investment-history margin, firms in both groups are closer to their most recent investment spike than the average. This is consistent with Lee (2022), who documents that recessions tend to follow periods of elevated lumpy investment by large firms. The relevant contrast is across recovery types: firms entering weak-recovery recessions are further from their last investment spike than firms entering non-weak-recovery recessions. Years since the last spike are 6.1 percent below the full-sample average before weak recoveries, compared

[2] For the Great Recession, we use 2007 because the NBER dates the recession as starting in December 2007, so annual firm observations mostly reflect pre-recession conditions.

[3] For this historical comparison, we use an annual Compustat–CRSP panel going back to the early 1960s. Because the exercise requires non-missing distance to default and enough investment history to measure the time since the last observed investment spike, the first usable pre-recession year is 1968.

with 18.5 percent below average before non-weak recoveries.

Taken in isolation, this real-side difference need not predict a weak recovery. In a standard lumpy-investment environment, a longer spell since the last major adjustment may indicate a larger accumulated capital gap and a higher likelihood of subsequent adjustment. The distinguishing feature of weak-recovery episodes is that this relatively delayed adjustment coincides with substantially weaker balance sheets. Before weak recoveries, distance to default is 6.5 percent below its full-sample average; before non-weak recoveries, it is 18.0 percent above average. Firms therefore enter weak-recovery downturns with materially more fragile financial positions. Weak recoveries are thus preceded by a joint real-financial state: firms are relatively further from their last major adjustment and, at the same time, financially less able to finance the next one.

2.2 Fact 2: Lumpier firms are financially more fragile

We next show that the real-financial connection documented around recessions is also present in the firm cross-section. We measure investment lumpiness as the concentration of a firm's investment over time.[4] Our main measure is the within-firm Herfindahl-Hirschman Index,

$$\text{HHI}_j = \sum_{t=1}^{T_j} \left(\frac{i_{jt}}{\sum_{s=1}^{T_j} i_{js}} \right)^2, \quad \text{where } i_{jt} > 0. \tag{1}$$

where i_{jt} is firm j's investment rate in year t, and T_j is the number of years in which firm j is observed with positive investment. The index is low when investment is spread smoothly over time and high when investment is concentrated in a few large episodes.

We collapse the panel to a firm-level cross section and estimate

$$\text{HHI}_j = \alpha + \sum_{s=1}^{S} \gamma_s \mathbf{1}[j \in s] + \beta \text{Average Financial Position}_j + \mathbf{\Gamma} X_j + \epsilon_j, \tag{2}$$

where γ_s denotes a major-industry fixed effect, and X_j includes average real sales growth, average real total assets, the average share of current assets, average firm age, and the average five-year standard deviation of real sales growth. The coefficient β summarizes the conditional association between a firm's average financial position and its investment lumpiness.

Figure 1 visualizes the cross-sectional relationship in Equation (2), using firm-level

[4] For this exercise, we restrict attention to the 1974–2024 period, when Compustat–CRSP coverage is broader and financing variables are more stable. We also require firms to have at least ten years of data, so that investment concentration is measured over sufficiently long within-firm histories.

Figure 1: Investment concentration and financial position

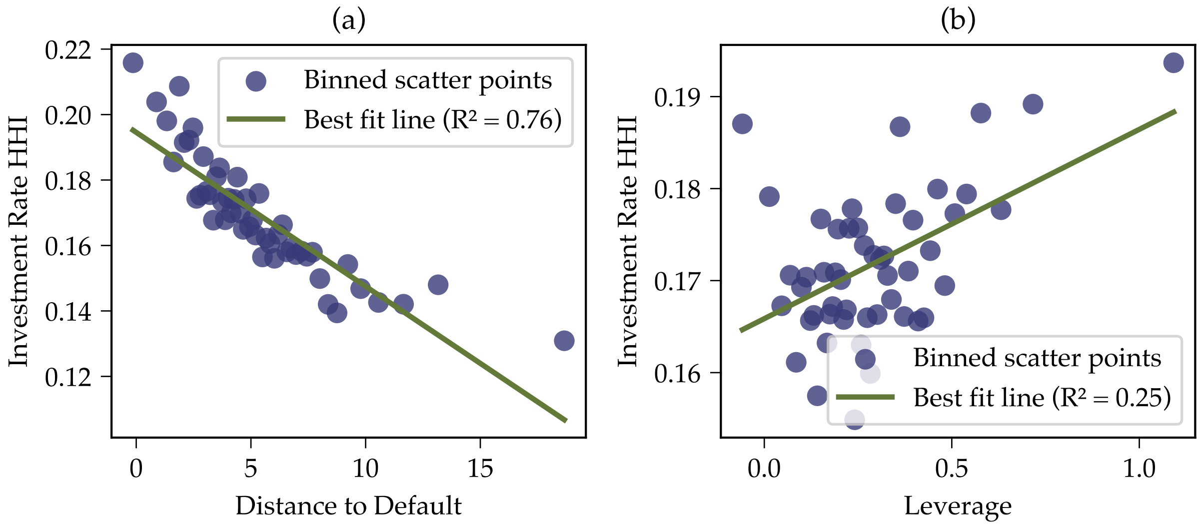

Note: Binned scatterplots of firm-level investment concentration and financial position. Panel (a) plots residualized investment HHI against distance to default; panel (b) plots residualized investment HHI against leverage. Each panel reports 50 firm-level bins and the fitted relationship from Equation (2).

bins together with the fitted conditional association summarized by β. The figure shows that lumpier firms are financially more fragile. Investment concentration is negatively related to distance to default and positively related to leverage: firms with lumpier investment histories operate closer to their default boundary and have higher leverage. Since lower distance to default indicates greater default risk and a higher implied external finance premium, while higher leverage indicates greater reliance on debt finance, both patterns point to weaker balance sheets among firms with more concentrated investment.[5]

We interpret this evidence as descriptive rather than causal. Persistent firm characteristics, growth opportunities, sectoral cyclicality, or financing needs may jointly shape both investment timing and financial fragility. The relevant takeaway is therefore not one of identification, but of comovement: in the firm cross-section, concentrated real adjustment is systematically associated with weaker balance sheets.[6]

[5]Distance to default is widely used as a market-based proxy for credit risk and external finance premia; see Gilchrist & Zakrajšek (2012) and Ottonello & Winberry (2020).

[6]Note also that this pattern is robust to alternative measurement choices. The association remains when we use other proxies for financial position, including liquidity and interest expenses, and when we replace the HHI with alternative measures of investment lumpiness. The full set of results is reported in Table A2 in Section A.3.

Figure 2: Percent of firms with investment and net debt issuance spikes over time

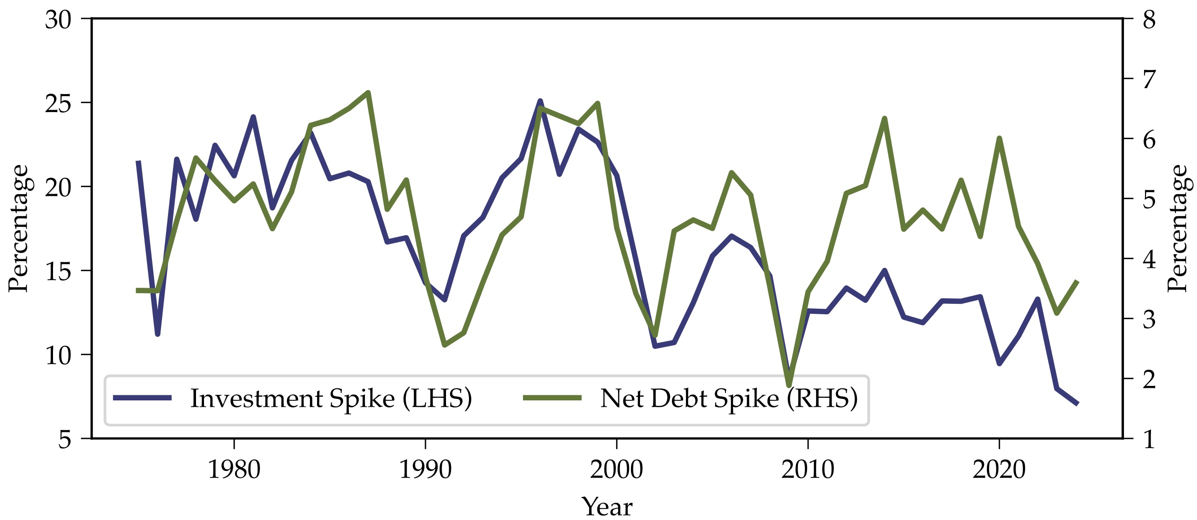

Note: This figure plots the percent of firms with an investment spike and the percent of firms with a net debt issuance spike between 1975 and 2024. The left axis measures the percent of firms with an investment spike, while the right axis measures the percent of firms with a net debt issuance spike.

2.3 Fact 3: Investment and debt spikes co-move

The previous fact shows that investment lumpiness and financial fragility are linked in the firm cross-section. We now show that the same real-financial connection is visible in the aggregate time series. Using the same 1974–2024 sample, we compare the share of firms with investment spikes to the share with large net debt issuance episodes. We define a debt spike as net debt issuance greater than or equal to 20 percent of total liabilities.

Figure 2 shows that investment and debt issuance spikes move closely together over time. Years with a larger share of firms undertaking large investment episodes are also years with a larger share of firms issuing substantial amounts of debt. This co-movement suggests that large capital adjustments are often financed externally. Financial conditions may therefore shape not only the level of investment, but also the timing of large adjustment episodes.[7] In the Appendix we also document that, at the firm level, large investment spikes are associated with an increase in debt and net debt issuance rate, with a larger probability of a debt issuance spike and a decline in liquidity.

Taken together, the three stylized facts point to a real-financial mechanism in recovery dynamics. Weak-recovery recessions are preceded by relatively less recent investment spikes and weaker balance sheets. In the firm cross-section, lumpier investment histories

[7] This pattern is not driven by the particular threshold used to define spikes. In Appendix A.4, we vary the spike cutoff between 10 and 30 percent and show that the time-series correlation between investment and debt issuance spikes remains robust. These results are reported in Figure A2.

are associated with greater financial fragility. In the aggregate time series, large investment episodes coincide with large debt issuance episodes. These patterns suggest that the timing of investment spikes depends not only on firms' desired capital adjustment, but also on their ability to finance that adjustment. The model below formalizes this interaction by combining non-convex capital adjustment with financial frictions.

3 Theory and Illustration

This section develops a simple firm investment model to rationalize the close link between investment lumpiness and financial conditions documented in the previous section. The model establishes that financial frictions alone can make firms choose non-adjustment rather than adjustment. We further demonstrate how these frictions amplify real adjustment costs, producing what we term *lumpy investment spirals*. To highlight the mechanism as transparently as possible, we deliberately adopt a parsimonious framework that nests financial frictions in their simplest form. Specifically, we assume that external financing is more costly than internal financing, thereby introducing a wedge that captures the essence of financial constraints without requiring additional structure on the source of the friction.

3.1 Model Setup

Consider a representative firm that enters the first period with capital k and productivity z. Capital depreciates at rate δ. Output is produced using a Cobb–Douglas technology with decreasing returns to scale:

$$\pi(k, z) = zk^{\alpha}. \tag{3}$$

Internal resources are current profits plus undepreciated capital:

$$\psi(k, z) = \pi(k, z) + (1 - \delta)k. \tag{4}$$

The firm uses its internal resources and, if needed, external finance to choose next-period capital for production in the second period. There is no uncertainty, and so productivity tomorrow z' is known at the time of investment. After the second period the firm exits and distributes net worth as dividends. The firm chooses next-period capital to solve

$$\max_{k' \geq (1-\delta)k} \quad \psi(k, z) - \mathcal{R}(k', k, z)k' + \beta\psi(k', z'). \tag{5}$$

The object $\mathcal{R}$ is the average unit cost of financing and is defined as

$$\mathcal{R}(k', k, z) = \begin{cases} 1 & \text{if } k' \leq \psi(k, z), \\ \frac{\psi(k,z)}{k'} + \mathcal{R}^{ext}(k', k, z)\left(1 - \frac{\psi(k,z)}{k'}\right) & \text{if } k' > \psi(k, z). \end{cases} \tag{6}$$

Expression (6) implies that, if financing needs can be fully covered with internal resources, the average cost equals one. However, if the firm chooses k' such that external financing is required, the average cost rises with the share of external funding. The external financing cost schedule is parameterized as

$$\mathcal{R}^{ext}(k', k, z) = e_0 + e_1\left(k' - \psi(k, z)\right), \qquad e_0 \geq 1, \quad e_1 \geq 0, \tag{7}$$

where e_0 is the baseline external unit financing cost, so that $e_0 - 1$ is the baseline spread over internal finance, and e_1 captures the sensitivity of financing costs to external financing needs. This sensitivity could be interpreted as higher dependence on external financing leading to an increase in risk, similar to the micro-foundation provided in Ottonello & Winberry (2020).[8] $\mathcal{R}^{ext}$ overall can be interpreted as the cost a firm pays on external financing, which could be in the form of loans, bonds, or equity issuance.[9]

Finally, the specific form of $\mathcal{R}^{ext}$ nests the equity issuance cost of Hennessy & Whited (2007).

Remark 1 (External finance costs equivalent to equity financing cost)**.**
$\mathcal{R}^{ext}$ maps into the equity issuance costs in Hennessy & Whited (2007), with $e_0 - 1$ and e_1 corresponding to λ_1 and λ_2 in Hennessy & Whited (2007).

Relative to Hennessy & Whited (2007), the key difference is that we impose this structure directly on the external unit financing cost. This induces a discontinuity in the marginal cost of funds at the internal-finance boundary, while preserving concavity away from the kink. Moreover, e_0 and e_1 could be directly estimated in loan-level data, where e_0 would be interpreted as a gross external finance cost and $e_0 - 1$ as a spread.

[8] The external finance structure is consistent with micro-foundations such as Ottonello & Winberry (2020). Proposition A.2 and Remark 1 in Appendix B.1 show how $e_0 - 1$ and e_1 can be mapped to a micro-founded spread.

[9] The baseline model most closely resembles equity issuance. Proposition A.1 in Appendix B.1 shows how the formulation maps into a debt financing case with a small scale adjustment to $\mathcal{R}^{ext}$.

3.2 External financing costs can generate non-adjustment

We now proceed to illustrate how the financial friction can lead firms to choose non-adjustment. We classify a firm as undertaking an investment adjustment if $k' > (1 - \delta + \nu)k$, and as choosing non-adjustment otherwise. Our finding is summarized in the next proposition:

Proposition 1 (Financial frictions generate non-adjustment)**.**
For any $e_0 > 1$, the marginal cost curve is discontinuous at the internal-finance boundary. This discontinuity generates non-adjustment if

$$1 \leq \beta \left[\alpha z' \psi(k,z)^{\alpha-1} + 1 - \delta \right] \leq e_0, \qquad \textit{and} \qquad \pi(k,z) \leq \nu k. \tag{8}$$

Proof.
The firm's value function is given by

$$J(k,z) = \max_{k' \geq (1-\delta)k} \quad \psi(k,z) - \mathcal{R}(k',k,z)k' + \beta\psi(k',z'). \tag{9}$$

Taking the derivative with respect to k', the marginal benefit is

$$MB(k') = \beta \left[\alpha z'(k')^{\alpha-1} + 1 - \delta \right]. \tag{10}$$

The marginal cost is one below the internal-finance boundary. Above the boundary, total financing cost is

$$\psi(k,z) + \left[e_0 + e_1(k' - \psi(k,z)) \right] \left(k' - \psi(k,z) \right), \tag{11}$$

and therefore marginal cost is $e_0 + 2e_1(k' - \psi(k,z))$. Thus, the marginal cost function exhibits a discontinuity at $k' = \psi(k,z)$:

$$\lim_{k' \to (\psi(k,z))^-} MC = 1, \qquad \lim_{k' \to (\psi(k,z))^+} MC = e_0. \tag{12}$$

The lower bound in (8) implies that the firm wants to use all internal resources, while the upper bound implies that it does not want to cross into external finance. Hence the firm stops at the kink, $k' = \psi(k,z)$. Finally, since $\psi(k,z) = (1-\delta)k + \pi(k,z)$, the condition $\pi(k,z) \leq \nu k$ implies

$$k' = \psi(k,z) \leq (1 - \delta + \nu)k, \tag{13}$$

so the firm remains inside the non-adjustment region. □

An illustration of Proposition 1 is provided in Figure 3a. The figure depicts the case in

Figure 3: Marginal benefit and marginal cost of investment

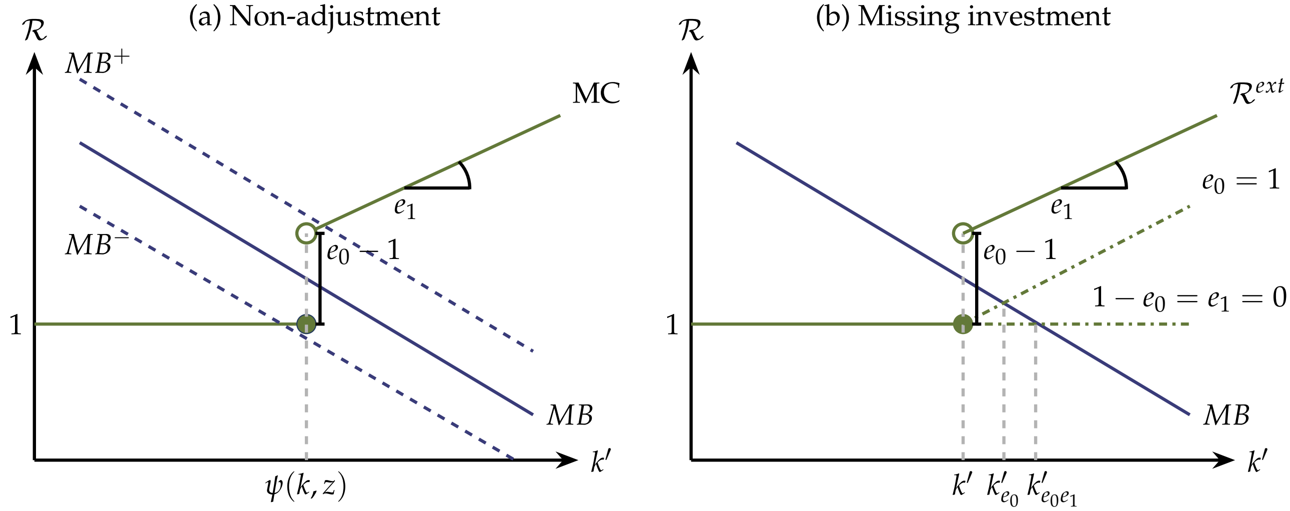

Note: This figure plots both the marginal cost and marginal benefit curves as a function of capital. The jump in the marginal cost curve occurs when the firm exhausts internal resources and further investment requires external finance. The left panel illustrates the non-adjustment region caused by a discontinuity in the marginal cost curve. The right panel illustrates the missing investment caused by e_0 and e_1.

which the marginal benefit at the internal-finance boundary lies within the discontinuity of the marginal cost schedule. In this situation, the firm uses its internal resources up to $\psi(k,z)$ but refrains from seeking external finance, since the marginal benefit does not justify paying the external financing wedge. Even small shifts in the marginal benefit curve within this range do not translate into changes in investment. Whether the firm adjusts depends on profitability relative to the non-adjustment band: if $\pi(k,z) > \nu k$, the firm can still move capital outside the non-adjustment band using internal resources; if $\pi(k,z) \leq \nu k$, stopping at internal resources leaves the firm in the non-adjustment region.

If, instead, the marginal benefit curve intersects the schedule above $k' = \psi(k,z)$, as indicated by the top dashed line, the firm combines internal funds with external financing in the amount $k' - \psi(k,z)$ and pays the external rate $\mathcal{R}^{ext}$. Alternatively, if the marginal benefit curve intersects the marginal cost schedule below $k' = \psi(k,z)$, as indicated by the bottom dashed line, internal funds suffice to finance the optimal investment.

Figure 3b illustrates the central role of e_0 and e_1 in shaping both the extent of non-adjustment and the aggregate investment shortfall generated by financial frictions. The fraction of firms that do not adjust, as well as the length of the non-adjustment region, naturally depends on the size of the spread $e_0 - 1$. When $e_0 = 1$, the marginal cost curve is continuous with only a kink, and the discontinuity does not induce non-adjustment. Conversely, the larger is e_0, the stronger the discontinuity in the marginal cost curve and the greater the fraction of firms that do not adjust. Since non-adjustment occurs only when internal resources are insufficient to finance desired capital accumulation, it predominantly affects less productive firms.

The same figure also illustrates the investment shortfall induced by these frictions. When $e_0 = 1$, the marginal cost curve exhibits no discontinuity, and tomorrow's capital rises from k' to k'_{e_0}. The gap between these two levels represents the portion of capital lost due to the discontinuity. If, in addition, e_1 is set to zero, capital increases further to $k'_{e_0 e_1}$, capturing the additional shortfall in capital attributable to the slope of the external financing cost schedule.

3.3 Lumpy investment spirals

The previous subsection establishes that a discontinuity in the marginal cost curve at the point where the firm accesses external financing can induce non-adjustment. We now show that, besides directly generating non-adjustment, financial frictions also amplify the effects of a real friction in the form of a fixed capital adjustment cost. Specifically, suppose that if the firm adjusts capital above the non-adjustment band $(1-\delta+\nu)k$, it pays a fixed adjustment cost ξ. Total financing needs are now

$$q(k',k) \equiv k' + \xi\, \mathbb{1}_{\{k'>(1-\delta+\nu)k\}}, \tag{14}$$

The financing schedule is the same as in (6) and (7), with financing needs $q(k',k)$ rather than k'. If the firm adjusts, its value is

$$J^*(k,z,\xi;e_0,e_1) = \max_{k'>(1-\delta+\nu)k} \quad \psi(k,z) - \mathcal{R}(q(k',k),k,z)q(k',k) + \beta\psi(k',z'), \tag{15}$$

whereas under non-adjustment, its value is

$$J^C(k,z;e_0,e_1) = \max_{(1-\delta)k\leq k'\leq(1-\delta+\nu)k} \quad \psi(k,z) - \mathcal{R}(k',k,z)k' + \beta\psi(k',z'). \tag{16}$$

The firm pays the adjustment cost if the full-adjustment value J^* is higher than the constrained value J^C. Equivalently, for given fundamentals, there is a maximum fixed cost that the firm is willing to pay in order to adjust. This adjustment-cost threshold is affected by $\mathcal{R}^{ext}$, as established in the following proposition:

Proposition 2 (External finance costs lower the adjustment-cost threshold)**.**

For given (k, z, z'), define

$$D(\xi; e_0, e_1) \equiv J^*(k, z, \xi; e_0, e_1) - J^C(k, z; e_0, e_1)$$

as the payoff gain from adjusting. Let ξ^* denote the indifference adjustment-cost threshold, implicitly defined by

$$D(\xi^*; e_0, e_1) = 0.$$

Then the firm adjusts iff $\xi \le \xi^*$. At thresholds where D is locally differentiable in $(\xi; e_0, e_1)$, higher external-finance costs lower this threshold:

$$\frac{\partial \xi^*}{\partial e_i} \le 0, \qquad i \in \{0, 1\}.$$

The inequality is strict whenever the adjusting firm uses external finance.

Proof.
First, D is strictly decreasing in ξ. The constrained value J^C does not depend on ξ. Under adjustment, total financing needs are

$$q(k', k) = k' + \xi.$$

For any fixed k', increasing ξ raises total financing needs and lowers the adjustment payoff. By the envelope theorem,

$$\frac{\partial J^*}{\partial \xi} = -MC(q(k'_A, k), k, z) < 0,$$

where MC is the marginal cost of financing total outlays. Below the external-finance boundary, $MC = 1$. Above the boundary,

$$MC = e_0 + 2e_1 \left(q(k'_A, k) - \psi(k, z)\right) > 0.$$

Therefore,

$$\frac{\partial D}{\partial \xi} = \frac{\partial J^*}{\partial \xi} < 0.$$

Thus D crosses zero at most once, so the adjustment decision is characterized by the cutoff ξ^*. Let

$$b_A(\xi) \equiv \left[q(k'_A(\xi), k) - \psi(k, z)\right]_+, \qquad b_{NA} \equiv \left[k'_{NA} - \psi(k, z)\right]_+$$

denote external financing needs under adjustment and non-adjustment, evaluated at their respective optimal choices. Adjustment requires

$$k'_A(\xi) > (1-\delta+\nu)k,$$

whereas non-adjustment requires

$$k'_{NA} \le (1-\delta+\nu)k.$$

Since $q(k'_A(\xi), k) = k'_A(\xi) + \xi$ and $\xi \ge 0$, we have

$$q(k'_A(\xi), k) > k'_{NA}.$$

Therefore,

$$b_A(\xi) \ge b_{NA},$$

with strict inequality whenever the adjustment option uses external finance. By the envelope theorem,

$$\frac{\partial J^*}{\partial e_0} = -b_A, \qquad \frac{\partial J^C}{\partial e_0} = -b_{NA}, \qquad \frac{\partial J^*}{\partial e_1} = -b_A^2, \qquad \frac{\partial J^C}{\partial e_1} = -b_{NA}^2.$$

Hence

$$\frac{\partial D}{\partial e_i} = \frac{\partial J^*}{\partial e_i} - \frac{\partial J^C}{\partial e_i} \le 0, \qquad i \in \{0,1\},$$

with strict inequality whenever the marginal adjusting firm uses external finance. At the indifference threshold,

$$D(\xi^*; e_0, e_1) = 0.$$

At thresholds where D is locally differentiable in (ξ, e_0, e_1), and since $\partial D/\partial \xi < 0$, the implicit-function argument gives

$$\frac{\partial \xi^*}{\partial e_i} = -\frac{\partial D/\partial e_i}{\partial D/\partial \xi}.$$

Since $\partial D/\partial e_i \le 0$ and $\partial D/\partial \xi < 0$, it follows that

$$\frac{\partial \xi^*}{\partial e_i} \le 0.$$

The inequality is strict whenever the marginal adjusting firm uses external finance. □

Proposition 2 establishes how external financing costs shape the firm's decision to pay the fixed adjustment cost. Higher e_0 and e_1 reduce the largest adjustment cost that the firm is willing to absorb, thereby shrinking the set of fixed-cost draws under which adjustment is optimal. The two parameters operate through related but distinct channels. The fixed component e_0 lowers the value of any adjustment that requires external finance, while the slope parameter e_1 penalizes larger financing gaps and therefore also reduces desired capital on the intensive margin. Adjustment is not pinned down by the fixed cost alone, however. For a given ξ, productivity also shifts the payoff gain from adjustment. Appendix B.3 provides the analogous productivity-threshold characterization, showing that higher external financing costs raise the productivity level required to justify adjustment.

We now turn to the second step of the spiral: when a firm chooses non-adjustment, capital depreciation lowers internal resources, increasing the external-finance gap at the next adjustment. This drift in internal resources amplifies the effects of external financing costs:

Proposition 3 (Lower internal resources raise external financing costs)**.**
Suppose the firm adjusts capital, remains in the external-finance region, and the adjusting optimum is interior. Let $\omega \geq 0$ denote a drift down in internal resources, so that

$$\psi_\omega(k,z) \equiv \psi(k,z) - \omega.$$

Let $k'_A(\omega)$ denote the optimal adjusting choice. An increase in ω raises the financing gap:

$$\frac{\partial}{\partial \omega}\left(q(k'_A(\omega),k) - \psi_\omega(k,z)\right) \geq 0.$$

Therefore, the external financing cost is weakly increasing in the drift down in internal resources:

$$\frac{\partial}{\partial \omega}\mathcal{R}^{ext} \geq 0,$$

with strict inequality when $e_1 > 0$.

Proof.
Let

$$b_A(\omega) \equiv q(k'_A(\omega),k) - \psi_\omega(k,z)$$

denote the external-finance gap of the adjusting firm. Conditional on adjustment, $q(k'_A(\omega),k) = k'_A(\omega) + \xi$. In the smooth external-finance region, the first-order condition is

$$\beta\left[\alpha z'(k'_A)^{\alpha-1} + 1 - \delta\right] = e_0 + 2e_1 b_A(\omega).$$

Define

$$M(k') \equiv \beta \left[\alpha z' (k')^{\alpha-1} + 1 - \delta \right].$$

Since $\alpha < 1$, $M'(k') < 0$. Differentiating the first-order condition with respect to ω gives

$$M'(k'_A) \frac{\partial k'_A}{\partial \omega} = 2e_1 \left(\frac{\partial k'_A}{\partial \omega} + 1 \right),$$

because

$$b_A(\omega) = k'_A(\omega) + \xi - \psi(k, z) + \omega.$$

Therefore,

$$\frac{\partial k'_A}{\partial \omega} = \frac{2e_1}{M'(k'_A) - 2e_1} \leq 0.$$

The firm weakly reduces desired capital when internal resources drift down. However, this reduction does not fully offset the mechanical increase in the financing gap:

$$\frac{\partial b_A}{\partial \omega} = 1 + \frac{\partial k'_A}{\partial \omega} = \frac{M'(k'_A)}{M'(k'_A) - 2e_1} \geq 0.$$

Hence lower internal resources raise the financing gap. Since

$$\mathcal{R}^{ext} = e_0 + e_1 b_A(\omega),$$

we have

$$\frac{\partial \mathcal{R}^{ext}}{\partial \omega} = e_1 \frac{\partial b_A}{\partial \omega} \geq 0,$$

with strict inequality when $e_1 > 0$. □

Proposition 3 establishes that lower internal resources raise the external-finance gap and, through e_1, raise the marginal cost of external finance. Importantly, this result allows the optimal adjusting capital choice to respond to the fall in internal resources. Although the firm reduces desired capital, it does not reduce it one-for-one, so the financing gap still rises.

The comparative static in Proposition 3 can be interpreted as the two-period representation of a dynamic non-adjustment spell. To see this, consider a sequence of two-period adjustment problems indexed by the number of periods since the firm's last large adjustment. If the firm last adjusted to capital k^0 and does not adjust for n periods, inherited capital evolves as

$$k^n = (1-\delta)^n k^0,$$

and internal resources at the next adjustment opportunity are

$$\psi^n(k^0, z) = \pi(k^n, z) + (1 - \delta)k^n.$$

For fixed productivity, $\psi^n(k^0, z)$ declines with n. The drift term ω in Proposition 3 can therefore be written as

$$\omega_n \equiv \psi^0(k^0, z) - \psi^n(k^0, z).$$

Thus, longer non-adjustment spells map into larger declines in internal resources. Proposition 3 shows that this raises the external-finance gap and, through e_1, the external unit financing cost.

Corollary 1 (Longer non-adjustment spells raise external financing costs)**.** *Suppose longer non-adjustment affects the next adjustment problem only through inherited internal resources, so that an increase in n maps into a weakly larger ω_n. Then the external-finance gap is weakly increasing in n, and $\mathcal{R}^{ext}$ is weakly increasing in n, with strict inequality when $e_1 > 0$ and ω_n strictly increases.*

The feedback between Propositions 2 and 3 gives rise to what we term lumpy investment spirals. In a dynamic setting, the mechanism works as follows. If a firm chooses non-adjustment today, capital depreciates and internal resources drift down. At the next adjustment opportunity, lower internal resources increase external financing needs and raise $\mathcal{R}^{ext}$ (Proposition 3). Higher external financing costs then make adjustment harder by lowering the adjustment-cost threshold ξ^* (Proposition 2). If the firm again chooses non-adjustment, the same force repeats until the firm undertakes a sufficiently large adjustment and exits the spiral.

4 Structural Model

We now describe the quantitative model used to evaluate the aggregate importance of the lumpy investment spiral. Time is discrete and infinite. There is a fixed mass of heterogeneous firms $j \in [0, 1]$ that face idiosyncratic productivity risk, aggregate productivity risk, lumpy real adjustment frictions, and costly external finance. The household is a representative agent who consumes, works, and owns firms. We denote the aggregate state by

$$S \equiv \{\Phi, A, \bar{\xi}, e_0\},$$

where Φ is the distribution of firms over individual states, A is aggregate productivity, $\bar{\xi}$ is the upper bound of the fixed adjustment cost distribution, and e_0 is the intercept of the external finance cost schedule.

4.1 Firms

At the beginning of each period, a firm enters with idiosyncratic productivity z and predetermined capital stock k. The firm observes the aggregate state S and rationally expects the law of motion for aggregate allocations. The firm produces using capital and labor, receives current operating profits, draws a fixed adjustment cost, and then chooses whether to use the full adjustment technology or remain within a constrained adjustment set.

Production Each firm operates a decreasing-returns Cobb-Douglas technology,

$$y = Azk^{\alpha}n^{\gamma}, \tag{17}$$

with $\alpha > 0$, $\gamma > 0$, and $\alpha + \gamma < 1$. Idiosyncratic productivity follows a log-AR(1) process,

$$\log(z') = \rho_z \log(z) + \sigma_z \epsilon'_z, \qquad \epsilon'_z \sim N(0,1), \tag{18}$$

where $0 \leq \rho_z < 1$ and $\sigma_z > 0$. The aggregate economy is driven by three shocks: aggregate productivity, the upper bound of the fixed adjustment cost distribution, and the intercept of the external finance cost schedule. These states follow

$$\log(A') = \rho_A \log(A) + \sigma_A \epsilon'_A, \qquad \epsilon'_A \sim N(0,1), \tag{19}$$

$$\log\left(\frac{\bar{\xi}'}{\bar{\xi}_{ss}}\right) = \rho_\xi \log\left(\frac{\bar{\xi}}{\bar{\xi}_{ss}}\right) + \sigma_\xi \epsilon'_\xi, \qquad \epsilon'_\xi \sim N(0,1), \tag{20}$$

$$e'_0 - e_{0,ss} = \rho_e(e_0 - e_{0,ss}) + \sigma_e \epsilon'_e, \qquad \epsilon'_e \sim N(0,1), \tag{21}$$

with $0 \leq \rho_A, \rho_\xi, \rho_e < 1$ and $\sigma_A, \sigma_\xi, \sigma_e > 0$. The real-friction shock is multiplicative in $\bar{\xi}$, while the financial shock is additive in e_0. In the numerical implementation, the discretized aggregate-state grid keeps $\bar{\xi} > 0$ and $e_0 > 1$ in every aggregate state.

Given the wage $w(S)$, labor is chosen statically:

$$\max_n Azk^{\alpha}n^{\gamma} - w(S)n. \tag{22}$$

The optimal labor demand and output are

$$n(k,z,S) = \left(\frac{\gamma A z k^{\alpha}}{w(S)}\right)^{\frac{1}{1-\gamma}}, \tag{23}$$

$$y(k,z,S) = (Az)^{\frac{1}{1-\gamma}} \left(\frac{\gamma}{w(S)}\right)^{\frac{\gamma}{1-\gamma}} k^{\frac{\alpha}{1-\gamma}}. \tag{24}$$

Static profits, net of labor costs, are

$$\pi(k,z,S) = (1-\gamma)y(k,z,S). \tag{25}$$

Real constraints Firms can adjust capital using two technologies. If the firm pays a fixed adjustment cost, it can freely choose next-period capital. The fixed cost is a labor cost $\xi w(S)$, where

$$\xi \sim G(\cdot;\bar{\xi}), \qquad \xi \in [0,\bar{\xi}].$$

The aggregate state $\bar{\xi}$ shifts the support of the fixed-cost distribution over the cycle. If the firm does not pay the fixed adjustment cost, it can only choose next-period capital inside a constrained set. We write this set as

$$\Omega(k) = [(1-\delta-\nu)k, (1-\delta+\nu)k]. \tag{26}$$

Thus, non-adjustment does not require capital to be constant. It allows small replacement investment around depreciated capital, but rules out the large adjustment that requires paying the fixed cost.

4.2 Financial intermediaries

Firms can finance capital purchases using internal resources and, if needed, external finance. The internal resources available for investment are

$$\psi(k,z,S) = \pi(k,z,S) + \theta(1-\delta)k, \tag{27}$$

where θ governs the fraction of undepreciated capital that can be used as internal funds. For a candidate capital choice k', required external finance is

$$b(k',k,z,S) = k' - \psi(k,z,S). \tag{28}$$

External finance has gross cost

$$R^{ext}(b, S) = e_0 + e_1 b, \qquad e_0 > 1, \quad e_1 > 0. \tag{29}$$

The average financing cost associated with the candidate choice k' is

$$R(k', k, z, S) = \begin{cases} 1, & \text{if } k' \leq \psi, \\ \dfrac{\psi}{k'} + R^{ext}(k' - \psi, S)\left(1 - \dfrac{\psi}{k'}\right), & \text{if } k' > \psi, \end{cases} \tag{30}$$

where $\psi = \psi(k, z, S)$ in Equation (30). If internal funds cover the capital purchase, the average cost is one. If the firm requires external funds, the cost is a weighted average of internal finance and external finance, with the external share priced by the current aggregate financing condition e_0.

4.3 Firm's problem

Let $m(S, S')$ denote the household stochastic discount factor. The firm value is

$$J(k, z, S) = \pi(k, z, S) + \int_0^{\bar{\xi}} \max\left\{J^*(k, z, \xi, S), J^c(k, z, S)\right\} dG(\xi; \bar{\xi}). \tag{31}$$

The first term is current operating profit. After production, the firm draws ξ and chooses between paying the fixed cost and using the full adjustment technology, or avoiding the fixed cost and choosing capital within $\Omega(k)$.

Adjustment Conditional on paying the fixed adjustment cost, the firm obtains

$$J^*(k, z, \xi, S) = (1 - \delta)k - R^*(k, z, S)\xi w(S) + Q(k, z, S, R^*(k, z, S)). \tag{32}$$

The object $R^*(k, z, S)$ is the internally consistent financing cost faced by the adjusting firm. To define it, first fix a candidate financing cost $\widetilde{R}$. Taking $\widetilde{R}$ as given, the firm chooses next-period capital to solve

$$Q(k, z, S, \widetilde{R}) = \max_{k' \geq 0} \left\{-\widetilde{R}\, k' + \mathbb{E}\left[m(S, S')J(k', z', S')\right]\right\}. \tag{33}$$

Let

$$k^*(k, z, S, \widetilde{R}) \in \arg\max_{k' \geq 0} \left\{-\widetilde{R}\, k' + \mathbb{E}\left[m(S, S')J(k', z', S')\right]\right\} \tag{34}$$

denote the associated policy. The equilibrium financing cost is the fixed point

$$R^*(k,z,S) = R\left(k^*(k,z,S,R^*(k,z,S)),k,z,S\right). \tag{35}$$

Thus, the adjusting firm chooses capital taking a financing cost as given, and the financing cost is then required to be consistent with the external finance implied by that capital choice. The same cost $R^*(k,z,S)$ applies to the fixed adjustment cost $\xi w(S)$ because the firm must finance the total resources needed to execute the adjustment.

Non-adjustment If the firm does not pay the fixed adjustment cost, it chooses next-period capital within the constrained set:

$$J^c(k,z,S) = \max_{k^c \in \Omega(k)} \left\{(1-\delta)k - k^c + \mathbb{E}\left[m(S,S')J(k^c,z',S')\right]\right\}. \tag{36}$$

The non-adjustment option does not require paying the fixed adjustment cost and does not use the costly external-finance technology. This option captures small replacement investment around depreciated capital.

4.4 Household

We close the model with a representative household who consumes, saves, and supplies labor. The household specification follows Khan & Thomas (2008) and Bachmann et al. (2013). Preferences are

$$U(C,L) = \log(C) - \eta L, \tag{37}$$

where C is consumption, L is labor supply, and η governs the disutility of labor.

The recursive household problem is

$$V(a,S) = \max_{C,L,a'} \quad \log(C) - \eta L + \beta \mathbb{E} V(a',S') \tag{38}$$

$$\text{s.t.} \quad C + \int a'(S')\, d\mathcal{P}(S'|S) = a + w(S)L. \tag{39}$$

The household trades state-contingent claims to firm profits. The stochastic discount factor is

$$m(S,S') = \beta \frac{C(S)}{C(S')}. \tag{40}$$

The intratemporal labor condition is $w(S) = \eta C(S)$.

4.5 Equilibrium

We are now able to define a recursive competitive equilibrium in this economy.

Definition 1 (Recursive Equilibrium). *An equilibrium of this model is a set of functions $J(k,z,S)$, $k'(k,z,S)$, $k^c(k,z,S)$, $\Phi(k,z,S)$, $R^*(k,z,S)$, m, w, a, C and L that solve the firm, financial intermediary and household problem and also clear the market for labor, output and equity, as described by the following conditions.*

1. *Given prices and the stochastic discount factor, firms solve Equation (31) with the adjustment, non-adjustment, and financing-cost policies defined above.*

2. *The household solves its problem, the stochastic discount factor is given by Equation (40), and the wage satisfies $w(S) = \eta C(S)$.*

3. *The distribution of firms evolves according to the idiosyncratic productivity transition, the aggregate shock processes, and the firm decision rules. Let $\kappa(k,z,\xi,S)$ denote the next-period capital selected after the fixed-cost draw. Aggregate capital is $K = \int k \, d\Phi(k,z)$ and next-period capital is*

$$K' = \int \int \kappa(k,z,\xi,S) \, dG(\xi;\bar{\xi}) \, d\Phi(k,z).$$

4. *Aggregate investment is*

$$I = \int \int [\kappa(k,z,\xi,S) - (1-\delta)k] \, dG(\xi;\bar{\xi}) d\Phi(k,z).$$

   Output is $Y = \int y(k,z,S) \, d\Phi(k,z)$, and the goods market clears:

$$Y = C + I.$$

5. *The labor market clears:*

$$L = \int n(k,z,S) \, d\Phi(k,z) + \int \int \xi \, \mathbb{1}\{J^*(k,z,\xi,S) \geq J^c(k,z,S)\} dG(\xi;\bar{\xi}) d\Phi(k,z).$$

6. *The market for firm claims clears by Walras' law.*

4.6 Solving and calibrating the model

Solution Method We solve the model using the global sequence-space solution method developed in Lee (2026), the repeated transition method (RTM). Because the firm problem

Table 2: Calibrated parameters and aggregate shock processes

Panel A: Firm-level parameters		
Parameter	Symbol	Value
Lumpy fixed-cost upper bound	$\bar{\xi}$	0.0300
Investment replacement range	ν	0.0378
External finance baseline	e_0	1.0188
External finance slope	e_1	0.0044
Panel B: Aggregate shock processes		
Shock and loading	Grid points	(ρ, σ)
A: multiplicative on productivity	3	$(0.859, 0.014)$
e: additive on e_0	2	$(0.720, 0.0056)$
$\bar{\xi}$: multiplicative on $\bar{\xi}$	2	$(0.850, 0.700)$

features nonconvex adjustment and an endogenous financing wedge, aggregate fluctuations in the recursive competitive equilibrium are potentially nonlinear. Since our main focus is endogenous state dependence in the propagation of aggregate shocks, we adopt this global nonlinear method, which recovers state-by-state transitions without imposing a local approximation. The approach exploits the ergodicity of the recursive competitive equilibrium: with a sufficiently long simulated path for the exogenous states, the relevant combinations of endogenous and exogenous states are realized. Each period's conditional expectation of value functions is computed by combining value functions from periods with similar endogenous states but different exogenous realizations. The method does not rely on a parametric law of motion. Instead, it requires only a metric that captures similarity of endogenous states across time. Implementation details are provided in Appendix C.1.

Calibration Each model period corresponds to one year. Table C1 in Appendix C.2 reports the fixed parameters. Most of these parameters follow Khan & Thomas (2008). The discount factor β implies an average annual real interest rate of 2.3%. The production parameters α and γ imply a capital share of 25.6% and a labor share of 64.0%. The depreciation rate is 9.0%, and the labor disutility parameter is set so that households work roughly one third of available time. The internal-resource parameter is $\theta = 1$, so firms can use all undepreciated capital as internal resources. The idiosyncratic productivity process uses the persistence and volatility in Khan & Thomas (2008).

Table 2 reports the calibrated parameters and aggregate shock processes. Panel A con-

Table 3: Calibration fit

Moment	Data	Model	Source
Panel A: Stationary moments			
Spike rate, $i/k > 0.20$	0.144	0.107	Zwick & Mahon (2017)
Mean i/k	0.104	0.118	Zwick & Mahon (2017)
Positive investment, $i/k > 0$	0.856	0.931	Winberry (2021)
Std. i/k	0.330	0.284	Ottonello & Winberry (2020)
Leverage, D/K	0.340	0.329	Ottonello & Winberry (2020); Crouzet & Mehrotra (2020)
Labor hours	0.330	0.336	8 working hours per day
External-finance spread	0.0204	0.0210	Ferreira et al. (2024)
Panel B: Time-series moments for the real-friction shock			
Std. of cross-sectional spike rate	0.045	0.045	Zwick & Mahon (2017)
Corr. of spike rate and aggregate I/K	0.760	0.500	Zwick & Mahon (2017)

Note: The appendix describes the construction of the cross-sectional standard deviation of investment rates, the implied stock leverage measure, and the borrower-weighted external-finance spread.

tains the four firm-level parameters. The upper bound of the fixed adjustment cost, the replacement-investment range, and the two parameters of the external finance schedule are calibrated to match stationary investment and financing moments. Panel B reports the aggregate shock processes. The aggregate TFP process follows Khan & Thomas (2008). The financial shock is an additive process for e_0, with persistence and volatility chosen to match the annual autocorrelation and time-series standard deviation of firm-level external-finance spreads in Ferreira et al. (2024). The real-friction shock is a multiplicative process for $\bar{\xi}$, calibrated to the time-series behavior of the cross-sectional investment-spike rate in Zwick & Mahon (2017). The three aggregate shocks are mutually independent and are discretized on a $3 \times 2 \times 2 = 12$ point joint grid.

Table 3 reports the model fit. The stationary cross-sectional moments are computed from the stationary distribution and the adjustment and non-adjustment policies. The model generates an investment spike rate of 10.7%, compared with 14.4% in Zwick & Mahon (2017), a mean investment rate of 11.8%, compared with 10.4%, and a cross-sectional standard deviation of investment rates of 28.4%, compared with 33.0% in Ottonello & Winberry (2020). It also matches the leverage and labor targets closely. Since debt is not a state variable in the model, the leverage moment is constructed from the model's flow of external finance and an empirical debt amortization rate; Appendix C.2 gives the details.

The borrower-weighted external finance spread is 2.10%, close to the 2.04% data target from Ferreira et al. (2024). The stationary cross-sectional fit has a normalized RMSE of 0.14. The last two rows report the time-series moments used to discipline the real-friction shock: the model matches the volatility of the cross-sectional spike rate exactly, and generates a positive correlation between spike rates and aggregate investment. The correlation is lower than in the data because the aggregate shocks are independent; increasing the real-friction shock to match the spike-rate volatility introduces spike-rate variation that need not co-move with aggregate I/K.

5 Equilibrium Analysis

In this section, we investigate the quantitative implications of the lumpy investment spiral. The theory in Section 3 shows that non-adjustment lowers internal resources, raises the future external-finance gap, and makes subsequent adjustment harder. The quantitative model allows us to ask whether this mechanism is strong enough to matter in equilibrium. We first show that the spiral leaves a clear imprint on the stationary allocation of capital across firms. We then connect this cross-sectional distortion to the sufficient statistics of Baley & Blanco (2021), which summarize the cumulative response of average capital in a lumpy-adjustment environment. Finally, we show directly that economies in which the spiral is active recover more slowly from recessions.

5.1 Stationary equilibrium and misallocation

We begin by comparing stationary equilibria across four versions of the model. The baseline model contains both real and financial frictions and is therefore the environment in which the spiral can operate. The no-financial-friction economy keeps the lumpy adjustment cost but removes the external-finance wedge, shutting down the feedback from non-adjustment to future financing costs. The financial-friction-only economy removes the lumpy adjustment cost but keeps the external-finance wedge, so firms do not experience persistent non-adjustment spells. The frictionless economy removes both frictions. These comparisons isolate the distinctive footprint of the spiral: misallocation that appears when financing costs interact with lumpy adjustment.

Table 4 shows that the baseline economy, where the spiral is active, is the most misallocated. The correlation between firm productivity and capital is 0.609 in the baseline model, compared to 0.842 in the frictionless economy. A financial friction without persistent non-adjustment does little to this statistic: the financial-friction-only economy looks

Table 4: Stationary cross-sectional allocation across model variants

Model	corr($\log z, \log k$)	Mass high-z/low-k	Mass low-z/high-k	Y share high-z/high-k
Baseline	0.609	0.0054	0.0027	0.209
No financial friction	0.670	0.0045	0.0013	0.222
Financial friction only	0.831	0.0003	0.0004	0.268
Frictionless	0.842	0.0001	0.0002	0.256

Note: High-z/low-k firms are firms in the top productivity quartile and bottom capital quartile. Low-z/high-k firms are firms in the bottom productivity quartile and top capital quartile. High-z/high-k firms are firms in the top quartile of both productivity and capital.

very similar to the frictionless economy. A lumpy economy without the financial feedback is more distorted, but the baseline is substantially more misallocated. The key difference is the spiral: firms that delay adjustment enter future periods with lower internal resources, so adjustment becomes more costly precisely when capital is already drifting away from its desired level.

The same mechanism appears in the tails of the distribution. The mass of high-productivity firms with low capital is 0.0054 in the baseline economy, compared to 0.0045 in the no-financial-friction economy and only 0.0001 in the frictionless economy. These firms are the natural candidates for the spiral: they have high desired capital, but non-adjustment erodes their internal resources and increases their reliance on costly external finance. At the same time, the baseline has the lowest output share produced by high-productivity, high-capital firms. The spiral therefore reallocates output away from the most efficient firms.

This pattern is important for interpreting the quantitative experiment. The main result is not merely that adding frictions worsens the allocation. Rather, the stationary distribution displays the specific form of misallocation predicted by the spiral: productive firms become undercapitalized when delayed adjustment and external-finance costs reinforce one another. The financial friction plays little role by itself because, absent lumpy adjustment, firms do not spend long periods drifting away from their desired capital. Its aggregate force appears through the spiral.

5.2 Non-adjustment spells and sufficient statistics

The stationary distribution also reveals why the spiral matters for aggregate dynamics. Baley & Blanco (2021) show that the cumulative response of average capital to an aggre-

Table 5: Baley-Blanco sufficient statistics across stationary equilibria

Model	$\text{Var}[x]$	$E[a]$	$\text{Var}[a]$	$\text{Cov}[x,a]$
Baseline	0.110	5.32	21.45	−0.99
No financial friction	0.117	4.54	16.46	−0.87
Financial friction only	0.089	0.29	0.45	−0.05
Frictionless	0.115	0.08	0.11	−0.01

Note: The statistic x is defined as $\log(k/z)$ and a is the number of periods since the firm last adjusted. $\text{Var}[x]$ is computed from the stationary distribution. The remaining statistics are computed from a simulated panel of 10,000 firms for 2,000 periods under the stationary decision rules, with a 500-period burn-in.

gate productivity shock can be summarized by moments of the firm distribution. Let

$$x_i \equiv \log(k_i/z_i)$$

denote firm i's capital-to-productivity gap, and let a_i denote the number of periods since firm i last paid the lumpy adjustment cost. The variance of x_i captures dispersion in capital relative to productivity. The covariance between x_i and a_i captures whether capital gaps are concentrated among firms deep into their non-adjustment spells.

Table 5 shows that $\text{Var}[x]$ is not the main discriminator across model variants. The more informative statistic for the spiral is $\text{Cov}[x,a]$. The baseline has $\text{Cov}[x,a] = -0.99$, compared to -0.87 in the no-financial-friction economy. In the economies without lumpy adjustment frictions, the covariance is close to zero because firms rarely experience meaningful non-adjustment spells.

The source of this stronger covariance is informative. Using

$$\text{Cov}[x,a] = \text{Corr}[x,a]\sqrt{\text{Var}[x]}\sqrt{\text{Var}[a]},$$

the move from the no-financial-friction economy to the baseline is driven primarily by the age distribution. The variance of non-adjustment age rises from 16.46 to 21.45, while the variance of x is slightly lower in the baseline. Thus the spiral does not mainly operate by creating more dispersion in capital gaps per se. It operates by spreading firms further across non-adjustment spell lengths.

To connect this statistic to recovery speed, let x_i^* be the post-shock target and define

$$g_i \equiv x_i^* - x_i,$$

so $g_i > 0$ means undercapitalization. When the target is approximately common, or when

$\text{Cov}(x_i^*, a_i)$ is small,

$$\text{Cov}(g_i, a_i) = \text{Cov}(x_i^*, a_i) - \text{Cov}(x_i, a_i) \approx -\text{Cov}(x_i, a_i).$$

A more negative Baley–Blanco covariance therefore means that older non-adjusters carry larger capital gaps.

Define the effective gap-closing rate on the recovery sample,

$$\chi_i \equiv \mathbb{E}_t \left[\frac{g_{it} - g_{i,t+1}}{g_{it}} \right], \qquad g_{it} > 0.$$

This object combines the extensive margin, whether the firm adjusts, and the intensive margin, how much of the gap it closes conditional on adjustment. The average gap evolves as

$$\bar{g}_{t+1} = \bar{g}_t - \mathbb{E}[\chi_i g_i].$$

Proposition 4 (Gap-persistence decomposition)**.**
Let $\zeta \equiv \text{Cov}(g_i, a_i)/\text{Var}(a_i)$ and $\lambda \equiv -\text{Cov}(\chi_i, a_i)/\text{Var}(a_i)$ be the best-linear-projection slopes of g and χ on age. Then the one-period persistence of the aggregate gap satisfies

$$\rho_g \equiv \frac{\bar{g}_{t+1}}{\bar{g}_t} = \underbrace{(1 - \bar{\chi})}_{\textit{frequency}} + \underbrace{\frac{\lambda\, \zeta\, \text{Var}(a_i)}{\bar{g}_t}}_{\textit{age gradient} \times \textit{dispersion}} - \frac{\text{Cov}(u_i, v_i)}{\bar{g}_t}, \tag{41}$$

where u_i, v_i are the projection residuals of χ_i, g_i on a_i. Dropping the residual term, for given gradients λ, ζ the persistence ρ_g is increasing in $\text{Var}(a_i)$ whenever $\lambda\zeta > 0$.

Proof.
From $\bar{g}_{t+1} = \bar{g}_t - \mathbb{E}[\chi_i g_i]$ and $\mathbb{E}[\chi g] = \bar{\chi}\bar{g} + \text{Cov}(\chi, g)$,

$$\rho_g = \frac{\bar{g}_{t+1}}{\bar{g}_t} = (1 - \bar{\chi}) - \frac{\text{Cov}(\chi_i, g_i)}{\bar{g}_t}.$$

Write the projections $g_i = \bar{g} + \zeta(a_i - \bar{a}) + v_i$ and $\chi_i = \bar{\chi} - \lambda(a_i - \bar{a}) + u_i$, with $u_i, v_i \perp a_i$ by construction. Then $\text{Cov}(\chi_i, g_i) = -\lambda\zeta\, \text{Var}(a_i) + \text{Cov}(u_i, v_i)$. Substituting gives (41). With $\text{Cov}(u, v)$ negligible and λ, ζ fixed, $\partial\rho_g/\partial\text{Var}(a) = \lambda\zeta/\bar{g}_t$, which has the sign of $\lambda\zeta$. □

The signs of the two gradients have natural interpretations. The condition $\zeta > 0$ means that older firms carry larger capital gaps. The condition $\lambda > 0$ means that older firms close a smaller fraction of those gaps. When both conditions hold, a wider distribution of non-adjustment ages raises the persistence of the average capital gap.

Proposition 5 (The spiral flips the age gradient)**.**
Index a firm by its non-adjustment duration a; let internal resources at the next adjustment opportunity be $\psi^a = \psi^0 - \omega_a$ with ω_a non-decreasing in a. Write $\chi(a) = h(a)\,\phi_A(a) + (1 - h(a))\,\phi_{NA}(a)$, where $h(a) \in [0,1]$ is the adjustment probability, $\phi_A(a) \in (0,1]$ the expected fraction of the gap closed conditional on adjusting, and $\phi_{NA}(a) \leq 0$ the fraction closed under non-adjustment ($\phi_A > \phi_{NA}$).

(i) ***Firms in the spiral.*** *In the structural model with $e_1 > 0$, $\chi(a)$ is non-increasing in a, strictly decreasing whenever the adjusting firm uses external finance; hence $\lambda \geq 0$, strictly $\lambda > 0$.*

(ii) ***Frictionless-finance benchmark.*** *If $e_0 = e_1 = 0$ and, holding idiosyncratic productivity fixed, the non-adjustment-driven capital drift widens the gap with a, then $\chi(a)$ is non-decreasing in a and $\lambda \leq 0$.*

Proof.
(i) Under the calibration, non-adjusters remain in the non-adjustment band and finance entirely from internal resources, so $b_{NA} = 0$ and V^{NA} is independent of the external-finance wedge and of ω. The adjuster, in contrast, borrows $b_A > 0$, and by Proposition 3 a fall in internal resources ($\omega \uparrow$) strictly increases its external-finance gap and, through $e_1 > 0$, its marginal financing cost, so $\partial V^A/\partial\omega < 0$. Hence the adjustment gain $D(\xi; a) = V^A - V^{NA}$ satisfies $\partial D/\partial\omega < 0$. Since $\partial D/\partial\xi < 0$ (Proposition 2), the implicit-function theorem applied to $D(\xi^*; a) = 0$ gives $\xi^{*\prime}(\omega) = -(\partial D/\partial\omega)/(\partial D/\partial\xi) < 0$, so the adjustment probability $h(a) = \xi^*(a)/\bar{\xi}$ is decreasing in a (using ω_a non-decreasing in a; Corollary 1). Conditional on adjusting, the optimal target k'_A is non-increasing in ω (Proposition 3), so the firm resets to a lower level and $\phi_A(a)$ is non-increasing in a. Finally, continued depreciation widens a non-adjuster's gap, so $\phi_{NA}(a) \leq 0$ is non-increasing. As a rises, $h(a)$ and $\phi_A(a)$ fall and weight shifts from $\phi_A > 0$ toward $\phi_{NA} \leq 0$; therefore $\chi(a)$ is non-increasing in a, strictly when the adjuster borrows and $e_1 > 0$. The projection slope of χ on a is then ≤ 0, i.e. $\lambda \geq 0$ (strict).

(ii) With $e_0 = e_1 = 0$ the external-finance wedge is absent ($\mathcal{R}^{ext} = 1$): adjusters pay only the resource cost of capital and reset to the unconstrained optimum, so $\phi_A(a) = 1$ for all a and $\xi^*(a)$ depends on a only through the adjustment gain $D(0; a) = V^A - V^{NA}$. Holding z fixed, a non-adjuster's capital $k^a = (1-\delta)^a k^0$ falls with a, widening the gap to the concave unconstrained optimum, so $D(0; a)$ is non-decreasing in a; hence $\xi^*(a)$ and $h(a) = \xi^*(a)/\bar{\xi}$ are non-decreasing in a. With $\phi_A \equiv 1$ and h non-decreasing, $\chi(a)$ is non-decreasing in a, so $\lambda \leq 0$. □

Proposition 5 characterizes the force operating on firms in the spiral. It does not imply that the age gradient is positive for all firms in the economy. Outside the spiral, the standard lumpy-adjustment force may dominate: older firms are further from their target, closer to their adjustment trigger, and may close a larger fraction of their gap when they adjust. Whether the spiral force dominates in the aggregate, and whether λ is positive in the recovery sample, is therefore a quantitative question.

In the calibrated economy, the spiral force dominates on the recovery sample. Among undercapitalized firms, the projection slope of χ on age is $\lambda = 0.006$, with a 95% confidence interval $[0.002, 0.009]$, and the gap gradient is $\zeta = 0.024$. The variance of non-adjustment age rises by 30% under the spiral, from 16.5 to 21.5, and the Baley–Blanco covariance $|\text{Cov}(x, a)|$ rises from 0.87 to 0.99. With $\lambda > 0$ and $\zeta > 0$, Proposition 4 implies that the increase in $\text{Var}(a)$ raises the persistence of average capital gaps.

Taken together, the two propositions deliver a simple aggregate prediction: under the spiral, the projection slope λ is positive and the variance of non-adjustment durations rises, and by Proposition 4 their product is what slows aggregate recovery. The mechanism is intuitive: the financial friction reduces the adjustment probability of firms deepest into their non-adjustment spells, leaving the largest capital gaps in the slowest-adjusting tail of the firm distribution; the spiral's wider dispersion of non-adjustment durations then concentrates more of the aggregate gap in this slow-moving tail, lengthening the half-life of recovery. The next subsection verifies that this distributional force translates into slower aggregate capital recovery.

5.3 Recovery speed

The previous subsection links the stationary distribution to the persistence of capital gaps. We now ask whether this mechanism appears in the full dynamic equilibrium. This step is important because the sufficient-statistic argument is local and distributional, while aggregate capital also reflects equilibrium prices, moving targets, and the full cross-section of firms.

For each regime, we estimate the conditional law of motion

$$\log K_{t+1} = a + \rho_K \log K_t,$$

conditioning on periods in which aggregate TFP has returned to its middle state. The

Figure 4: Capital recovery speed across friction regimes

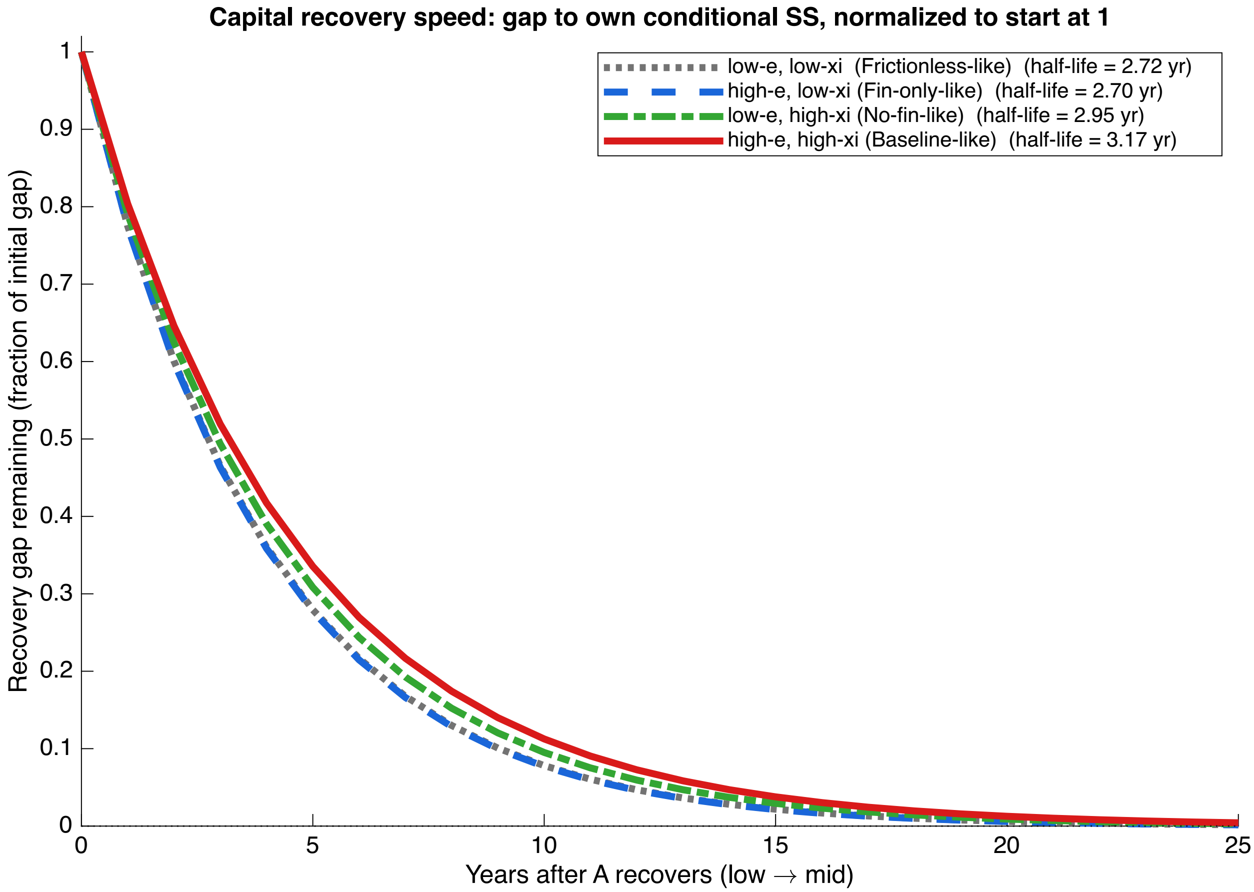

Note: The figure plots the fraction of the initial capital recovery gap that remains after aggregate TFP returns from the low state to its middle state. Each line normalizes the initial gap to one and measures convergence relative to the regime-specific conditional steady state.

conditional steady state is $K^* = \exp(a/(1-\rho_K))$, and the implied half-life is

$$\tau = \frac{\log(2)}{\log(1/\rho_K)}.$$

We then start each regime from the same recession trough, $K_0 = 0.92K_{ss}$, and compare the speed at which the capital gap closes.

Figure 4 and Table 6 show that the baseline-like regime recovers most slowly. The frictionless-like and financial-only-like regimes are nearly indistinguishable, with capital half-lives of 2.72 and 2.70 years. Thus, financial frictions alone do not materially slow aggregate recovery. When firms can adjust capital without persistent non-adjustment spells, the financial wedge has little aggregate force.

The real-only-like regime raises the half-life to 2.95 years, indicating that lumpy adjustment matters on its own. But the largest effect appears when lumpy adjustment in-

Table 6: Conditional capital recovery at middle aggregate productivity

Regime	ρ_K	Half-life
Low e, low $\bar{\xi}$ (frictionless-like)	0.7749	2.72
High e, low $\bar{\xi}$ (financial-only-like)	0.7738	2.70
Low e, high $\bar{\xi}$ (real-only-like)	0.7903	2.95
High e, high $\bar{\xi}$ (baseline-like)	0.8038	3.17

Note: The half-life is measured in years and is computed from the estimated conditional law of motion for aggregate capital at the middle TFP state.

teracts with external-finance costs. The baseline-like regime raises the half-life further to 3.17 years. Relative to the frictionless-like regime, the baseline-like economy recovers 17% more slowly.

This pattern is the quantitative signature of the spiral. Real frictions create non-adjustment spells. Financial frictions become powerful because they operate through those spells: delayed adjustment lowers internal resources, raises future external-finance needs, depresses the effective gap-closing rate of firms in the spiral, and spreads capital gaps across a wider tail of non-adjusters. The result is not simply that two frictions have larger effects than one. Rather, financial frictions have little effect in isolation but generate substantial persistence when they interact with lumpy adjustment. The recovery-speed exercise therefore confirms the mechanism identified by the sufficient statistics: the spiral slows aggregate recovery by placing capital shortfalls among firms that adjust late and close their gaps slowly.

6 Mechanism Validation

This section organizes the empirical evidence around the four links of the lumpy-investment spiral. First, investment spikes require external finance. Second, financing-intensive spikes weaken firms' balance sheets. Third, weaker balance sheets reduce the probability of future spikes, even conditional on the time elapsed since the last spike. Fourth, weak-recovery recessions are preceded by a firm distribution with more dispersed non-adjustment ages and a larger mass of firms that combine fragile balance sheets with accumulated investment pressure. We test these implications in annual Compustat data described in Section 2.

6.1 Investment spikes and financial conditions

The first two links in the empirical chain are that investment spikes require external finance and that financing-intensive spikes weaken firms' balance sheets. We begin with local projections around isolated investment spikes. To isolate sharp investment events rather than persistent investment episodes, we define an isolated spike as a spike that occurs after at least two years without a spike, and estimate

$$y_{i,t+h} = \beta_h \, \text{Spike}_{it}^{\text{iso}} + \rho_h y_{i,t-1} + \psi_{1h}\text{Spike}_{i,t-1}^{\text{iso}} + \psi_{2h}\text{Spike}_{i,t-2}^{\text{iso}} + \Gamma_h' X_{i,t-1} + \alpha_i + \delta_t + \varepsilon_{i,t+h}, \tag{42}$$

for horizons $h = -2, \ldots, 5$. The vector $X_{i,t-1}$ contains lagged firm controls: real sales growth, measured as the log change in real sales; firm size, measured as log real total assets; the share of current assets, measured as current assets divided by total assets; and cash flow, measured as operating income after depreciation divided by lagged gross property, plant and equipment. The sales-growth, size, and current-asset-share controls enter as standardized variables. The lagged dependent variable normalizes the coefficient in the period immediately before the event, while the lagged isolated-spike controls absorb the effects of recent isolated spikes. Firm fixed effects absorb time-invariant heterogeneity across firms, while year fixed effects absorb aggregate shocks. Standard errors are clustered by firm and year.

The local projections show that investment spikes are financing-intensive events. Appendix D reports the responses of several financing outcomes: net debt issuance increases, the probability of a debt-issuance spike rises, liquidity falls, and interest expenses increase around isolated investment spikes. These patterns are the first link in the mechanism. Large investment adjustments are not merely real reallocations inside the firm; they coincide with a discrete use of external finance and a drawdown of liquid resources.

The same events are followed by weaker financial positions. Figure 5 shows the response of distance to default. The coefficient one year before the spike is normalized to zero by the lagged dependent variable. After the spike, distance to default falls sharply: the coefficient is -0.54 one year after the investment episode and -0.26 two years after, with both estimates statistically significant. Thus, the firm-level event-study evidence connects the first two steps of the spiral: spikes are associated with external financing needs, and the balance-sheet position of the firm deteriorates after the spike. This evidence is descriptive, since investment decisions are jointly chosen with the firm's balance sheet, but it establishes that the relevant objects move in the direction implied by the model.

Figure 5: Distance to default after an isolated investment spike

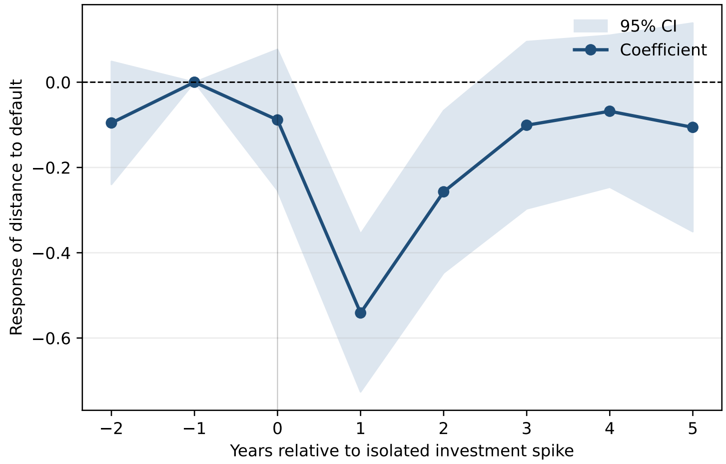

Notes: The figure reports local-projection coefficients from Equation (42). The shaded area is the 95 percent confidence interval. The event is an investment spike following at least two years without a spike. The specification includes one lag of the dependent variable.

To push the interpretation further, inspired by Zwick & Mahon (2017), we use tax policy as a source of variation in investment spikes. Federal bonus depreciation under section 168(k) allows firms to deduct a fraction of eligible investment immediately rather than depreciating the full cost gradually over the asset's tax life. In present-value terms, this lowers the tax price of new capital and creates an incentive to bring investment forward. The policy was introduced for investment placed in service after September 2001, with a 30 percent bonus rate, increased to 50 percent for part of 2003 and for 2004, and then expired over 2005–2007. It was reintroduced in 2008 at a 50 percent rate and remained in place in 2009; later extensions raised the rate further, but our baseline IV sample stops in 2011 to focus on the initial introduction and reintroduction episodes. The identifying variation is therefore a temporary fall in the user cost of capital, interacted with predetermined firm exposure. The appendix shows that the timing of bonus depreciation lines up with the spike margin with firms further into a non-adjustment spell more likely to spike around bonus-depreciation start and restart years, with no previous trends before the policy being implemented.

The exclusion restriction is that bonus depreciation affects distance to default primarily by inducing firms to undertake large investment episodes. The policy changes the timing of tax deductions for new investment, but it does not directly revalue existing debt, mechanically change the firm's pre-existing asset base, or alter default risk absent an investment response. If anything, the direct cash-tax effect conditional on investment should relax financing needs by increasing after-tax cash flow; that channel would tend to improve liquidity and distance to default. Finding that tax-induced spikes reduce distance to default is therefore consistent with the balance-sheet cost of lumpy investment dominating any direct tax-saving effect.

The baseline IV design uses a firm-level capital-intensity exposure inspired by the bonus-depreciation design in Zwick & Mahon (2017), but it does not reconstruct their exact asset-class depreciation-benefit measure. Compustat does not report, at the firm-year level, the tax asset-class composition needed to measure the share of investment that is eligible for bonus depreciation. We therefore use a simpler predetermined exposure: constructed capital K_{it}, using the perpetual inventory method, scaled by total book assets A_{it}. Firms with higher capital intensity should benefit more from the bonus depreciation and therefore react more to the policy. The firm exposure $\overline{K/A}_{i,1996\text{–}2000}$ is the average of this ratio over 1996–2000, after winsorizing K/A at the 1st and 99th percentiles in the full panel. The instrument interacts this predetermined exposure with the positive change in the statutory bonus-depreciation rate:

$$Z_{it}^{K/A} = \overline{K/A}_{i,1996\text{–}2000} \times \Delta^{+}\text{Bonus}_t.$$

The fixed firm component is predetermined relative to the 2001 and 2008 policy changes, while year fixed effects absorb the aggregate policy shock itself. We then estimate IV local projections,

$$Y_{i,t+h} = \beta_h \widehat{\text{Spike}}_{it} + \rho_h Y_{i,t-1} + \psi_{1h}\text{Spike}_{i,t-1} + \psi_{2h}\text{Spike}_{i,t-2} + \Gamma_h' X_{i,t-1} + \alpha_i + \delta_t + \varepsilon_{i,t+h}, \tag{43}$$

where $X_{i,t-1}$ contains lagged real sales growth, log real assets, the current-assets share of total assets, and cash flow, and $Y_{i,t+h}$ is distance to default. $\widehat{\text{Spike}}_{it}$ is the spike dummy instrumented with $Z_{it}^{K/A}$. The sample runs from 1990 to 2011.

Table 7 provides evidence for the second link in the mechanism: investment spikes induced by tax incentives weaken firms' financial positions on impact. The IV estimate is negative and statistically significant: a tax-induced investment spike is associated with an 8.81-unit decline in distance to default on impact. By contrast, the OLS estimate is also negative, but much smaller in magnitude and not statistically distinguishable from zero at conventional levels. This contrast is consistent with ordinary investment spikes being strongly endogenous: firms may choose to undertake large investments precisely when their expected prospects or financing capacity are relatively favorable, attenuating the adverse balance-sheet effect in OLS. The IV specification instead isolates variation in spikes generated by firms' exposure to positive bonus-depreciation changes. The first

Table 7: Tax-induced investment spikes and distance to default: impact estimates

	IV impact	OLS impact	First stage
Coefficient	-8.81	-0.10	0.20
	(3.49)	(0.07)	(0.07)
F-stat	-	-	8.53
N	42,176	42,176	42,176

Notes: IV and OLS impact columns report the coefficient on investment spikes in the horizon-zero distance-to-default specification with firm and year fixed effects. The first-stage column reports the coefficient on capital/assets exposure interacted with positive bonus-depreciation changes, its clustered standard error, and the corresponding F-statistic.

stage moves in the expected direction, with an estimated coefficient of 0.20, confirming that the tax instrument predicts investment spikes. The negative IV estimate supports the model's prediction that investment spikes, when induced by an external tax incentive rather than selected endogenously by firms, are followed by a deterioration in balance-sheet strength.

Capital-intensity exposure is not a randomized firm-level assignment and is coarser than the Zwick–Mahon industry-level present-value-of-depreciation measure. The fixed exposure component is predetermined and the aggregate tax change is statutory, but high-capital-intensity firms may differ from other firms in ways that matter for balance-sheet dynamics. The appendix therefore reports a separate robustness exercise based on lagged non-adjustment duration. The duration design shows that spike probabilities rise around bonus-depreciation starts primarily among firms that have been inactive for longer and that the policy induced spikes for these firms also decrease the distance to default, supporting the baseline instrument results.

6.2 Financial constraints and the probability of spiking

We next ask whether firms with weaker financial positions are less likely to undertake investment spikes. The baseline specification uses firm fixed effects, so the identifying variation comes from changes in a firm's financial position over time rather than permanent differences across firms. Because the likelihood of an investment spike depends on where a firm is in its investment cycle, we allow the association between financial constraints and spiking to vary with the time elapsed since the firm's last investment spike.

Figure 6: Investment-spike coefficients by distance to default

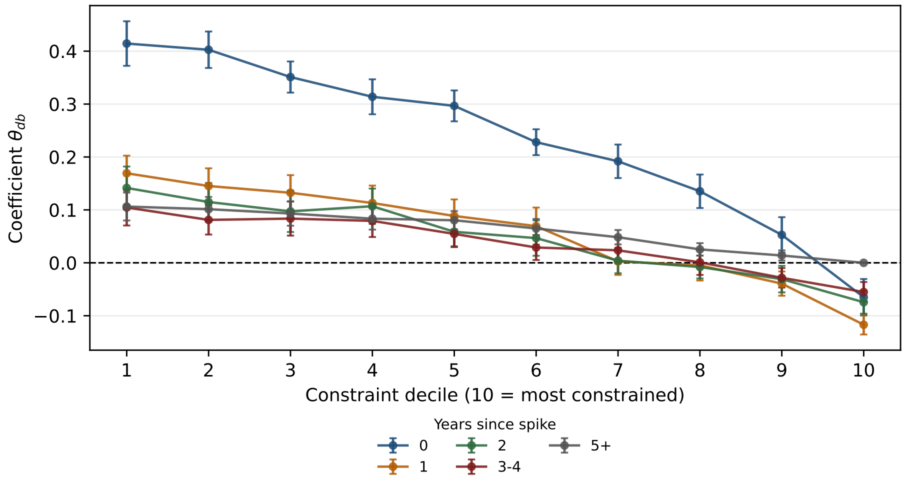

Notes: The figure plots the estimated coefficients $\widehat{\theta}_{db}$ from Equation (44) for each lagged distance-to-default decile and years-since-spike bin. Distance to default is oriented so that higher deciles indicate tighter financial conditions. The specification absorbs firm and year fixed effects and includes the lagged firm controls described in the text. Vertical bars report 95 percent confidence intervals.

We estimate linear probability models of the form

$$\text{Spike}_{it} = \sum_{d,b} \theta_{db} \mathbf{1}\{C \text{ decile}_{i,t-1} = d\} \mathbf{1}\{\text{Duration bin}_{i,t-1} = b\} + \Gamma' X_{i,t-1} + \alpha_i + \delta_t + u_{it}, \tag{44}$$

where C is distance to default, oriented so that higher deciles indicate tighter financial conditions. The duration bins are zero, one, two, three to four, and five or more years since the last spike. The vector $X_{i,t-1}$ is the same set of lagged firm controls used above: log real sales growth, log real total assets, current assets over total assets, and operating income after depreciation over lagged gross property, plant and equipment. Firm fixed effects α_i absorb time-invariant heterogeneity across firms and year fixed effects δ_t absorb aggregate shocks. The figure reports the estimated coefficients $\widehat{\theta}_{db}$ for each distance-to-default decile and duration-bin combination, together with 95 percent confidence intervals.

Figure 6 plots the estimated spike coefficients by distance-to-default decile within each time since last spike bin. Across duration bins, the coefficients tend to decline as firms become more financially constrained. Thus, conditional on a firm's position in its investment cycle, the same firm is less likely to undertake a spike in years in which it is closer to default. The pattern is particularly informative for the mechanism in the model: financial weakness is associated with delayed lumpy adjustment even after accounting for the fact that firms further into a non-adjustment spell may be closer to their next investment spike.

This graphical evidence is consistent with the specification using a standardized continuous constraint score. In that regression, the distance-to-default constraint score enters with a coefficient of -0.058, statistically different from zero. Thus, weaker balance sheets are associated with a lower probability of undertaking a large investment adjustment. In the Appendix, we repeat the interacted-coefficient analysis using alternative measures of firms' financial positions, including leverage, liquidity, and cash flow. These results deliver the same qualitative conclusion: tighter financial conditions are associated with less frequent investment spikes.

6.3 Pre-recession firm distributions

The previous subsections provide firm-level evidence on the feedback loop: investment spikes use external finance and weaken balance sheets, and weak balance sheets reduce the probability of a subsequent spike. The final step is to ask whether weak-recovery recessions are preceded by a firm distribution that makes aggregate spiral effects more likely. This diagnostic is measured before the recession starts, so it is less directly exposed to the concern that subsequent differences in investment dynamics simply reflect different aggregate shocks. We summarize two objects emphasized by the model: dispersion in non-adjustment age, a, and the mass of firms that combine weak balance sheets with a long time since the last investment spike. The data used are the same as in Section 2.1.

As established in Section 2.1, weak-recovery recessions are preceded, on average, by firms that are closer to default and further away from their last investment spikes. Here we test the model prediction more directly. The model goes beyond averages: recovery speed depends on how dispersed firms are across the non-adjustment-age distribution and on the mass of firms in the spiral state. To measure this state in the data, we classify firms as low-D2D if their distance to default is below the pooled 30th percentile, equal to 2.78. Long-since-spike firms are those at least three years since the last observed investment spike. Spiral mass is the share of firms satisfying both conditions.

Table 8 shows that weak-recovery recessions differ from historical non-weak recoveries along both margins. First, the variance of non-adjustment age is higher before weak recoveries: $\text{Var}(a)$ rises from 0.17 before non-weak recoveries to 0.72 before weak recoveries. Thus, weak recoveries are preceded not only by firms that are further into their in-

Table 8: Average firm distributions by recovery type

Recovery group	Var(a)	Low D2D (%)	Long since spike (%)	Spiral mass (%)	Cap.-weighted spiral (%)
Non-weak recoveries	0.17	18.2	39.0	6.4	3.8
Weak recoveries	0.72	35.2	38.6	11.9	11.6

Notes: The table averages recession-level moments separately for weak and historical non-weak recoveries. Weak recoveries are 1990–91, 2001, and 2007–09. Historical non-weak recoveries are 1969–70, 1973–75, and 1980–82. The 2020 pandemic recession is classified as a non-weak recovery but is excluded from the table because its shutdown-and-reopening dynamics make it a special episode relative to standard business-cycle recessions. Let $i_{it} \equiv \text{CAPX}_{it}/\text{PPEGT}_{i,t-1}$. For the age moment, the adjustment dummy is $\text{Adj}_{it}^{3.78} \equiv \mathbf{1}\{i_{it} \geq 0.0378\}$, and a is the number of years since the last observed year with $\text{Adj}_{it}^{3.78} = 1$. Low-D2D firms are those below the pooled 30th-percentile cutoff of distance to default, equal to 2.78. Long-since-spike firms are those at least three years since the last observed investment spike, where a spike is $i_{it} > 0.20$. Spiral mass is the share of firms satisfying both conditions. Capital-weighted spiral mass weights each firm by real capital in the pre-recession year before averaging across recessions.

vestment cycles on average, but also by a more dispersed distribution of non-adjustment ages. Appendix D.3 shows that this pattern is robust to using a stricter 1 percent threshold to define non-adjustment years.

Second, weak-recovery recessions have more firms in the joint real-financial region emphasized by the spiral. Before the 1990–91, 2001, and 2007–09 recessions, 11.9 percent of firms are both in the bottom 30 percent of distance to default and at least three years since the last spike, compared with 6.4 percent before historical non-weak episodes. The capital-weighted comparison is even sharper: spiral firms account for 11.6 percent of pre-recession capital before weak recoveries, compared with only 3.8 percent before historical non-weak recoveries. Since the overall share of firms with a long time since their last spike is similar across the two groups, this pattern is not simply about having more firms far from their last spike. It is about which of those firms enter the downturn with fragile balance sheets.

Taken together, the exercises validate the micro mechanism in the model and provide macro evidence in support of its aggregate effects. Investment spikes are followed by weaker financial conditions, weak financial conditions predict lower future spike probabilities, and a tax incentive that lowers the cost of capital raises spike probabilities and generates weaker balance sheets among the induced spikers. The pre-recession distribution adds a macro diagnostic: weak-recovery recessions are precisely those preceded by more dispersed non-adjustment ages and a larger mass of firms in the joint financial-real friction region. This pattern is the empirical analog of the model's spiral: lumpy investment strains the balance sheet, the weakened balance sheet makes future adjustment less likely, and aggregate recoveries are weaker when more firms enter the downturn with both accumulated adjustment pressure and fragile financial positions.

7 Conclusion

Why do recessions with similar initial declines recover at different speeds? This paper argues that the answer depends on the joint distribution of firms' financial positions and accumulated investment needs. When firms enter a downturn with weak internal resources and large capital gaps, the recovery is not governed only by how far capital is from its desired level. It also depends on whether the firms carrying those gaps can finance the adjustment required to close them. The interaction between financial fragility and lumpy investment therefore becomes a state variable for aggregate recovery dynamics.

We show that this interaction generates a lumpy-investment spiral. Fixed adjustment costs create periods in which firms do not adjust capital. Costly external finance changes the consequences of those delays. A firm that postpones adjustment allows capital to drift further from its desired level, while its internal resources deteriorate relative to the investment required to close the gap. The next adjustment then requires more external finance, making adjustment more costly and further reducing the incentive to adjust. The spiral therefore links the timing of real adjustment to the firm's balance-sheet position.

In the quantitative model, this mechanism changes both the level of misallocation and the speed of recovery. Financial frictions alone have limited aggregate effects. Their force appears when they interact with lumpy adjustment: they worsen the allocation of capital across firms, lengthen non-adjustment spells, and raise dispersion in the time since firms last adjusted. This distributional change is central for recovery speed. In a standard lumpy-adjustment economy, firms that have gone longer without adjusting tend to be closer to adjustment and help capital gaps close quickly. The spiral weakens this self-correcting force by making adjustment more expensive for firms deep into their spells. As a result, aggregate capital gaps close more slowly and recoveries become more persistent.

The empirical evidence supports the two sides of the mechanism and its aggregate implication. Large observed investment events are followed by weaker financial positions, and firms with weaker balance sheets are less likely to undertake subsequent large investment events. These patterns show that real adjustment strains the balance sheet and that financial fragility delays future adjustment. At the aggregate level, weak recoveries are preceded by a firm distribution with higher dispersion in non-adjustment age and a larger mass of firms that combine low distance to default with a long time since their last investment spike. This is the empirical counterpart of the distributional state that activates the spiral in the model.

The broader lesson is that financial conditions affect recoveries not only by changing the average cost of funds, but by reshaping which firms adjust and when. A downturn that hits an economy with many firms carrying both accumulated investment needs and fragile balance sheets can leave capital gaps unresolved for longer. Modeling the joint distribution of financial positions and investment histories is therefore central to understanding why some recessions rebound quickly while others remain persistently weak.

References

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

Bachmann, R. & Ma, L. (2016). Lumpy investment, lumpy inventories. *Journal of Money, Credit and Banking*, 48(5), 821–855.

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

Baley, I. & Blanco, A. (2024). *The Macroeconomics of Irreversibility*. FRB Atlanta Working Paper 2024-17, Federal Reserve Bank of Atlanta.

Beaudry, P., Galizia, D., & Portier, F. (2018). Reconciling hayek's and keynes' views of recessions. *The Review of Economic Studies*, 85(1), 119–156.

Begenau, J. & Salomao, J. (2018). Firm Financing over the Business Cycle. *The Review of Financial Studies*, 32(4), 1235–1274.

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

Beraja, M. & Wolf, C. K. (2022). Demand composition and the strength of recoveries.

Bernanke, B. S., Gertler, M., & Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of macroeconomics*, 1, 1341–1393.

Buera, F. & Karmakar, S. (2022). Real effects of financial distress: the role of heterogeneity. *The Economic Journal*, 132(644), 1309–1348.

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

Cloyne, J., Ferreira, C., Froemel, M., & Surico, P. (2023). Monetary Policy, Corporate Finance, and Investment. *Journal of the European Economic Association*. jvad009.

Cooper, R., Haltiwanger, J., & Power, L. (1999). Machine replacement and the business cycle: Lumps and bumps. *American Economic Review*, 89(4), 921–946.

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

Crouzet, N. (2018). Aggregate implications of corporate debt choices. *The Review of Economic Studies*, 85(3), 1635–1682.

Crouzet, N. & Mehrotra, N. R. (2020). Small and large firms over the business cycle. *American Economic Review*, 110(11), 3549–3601.

Donaldson, P. & Wieland, J. F. (2025). Why are some recoveries weak and others strong? Working paper, September 21, 2025.

Fernald, J., Inklaar, R., & Ruzic, D. (2025). The productivity slowdown in advanced economies: Common shocks or common trends? *Review of Income and Wealth*, 71(1), e12690.

Ferreira, T. R. T., Ostry, D. A., & Rogers, J. (2024). *Firm financial conditions and the transmission of monetary policy*. Bank of England working papers 1093, Bank of England.

Galí, J., Smets, F., & Wouters, R. (2012). Slow recoveries: A structural interpretation. *Journal of Money, Credit and Banking*, 44, 9–30.

Garga, V. & Singh, S. R. (2021). Output hysteresis and optimal monetary policy. *Journal of Monetary Economics*, 117, 871–886.

Gilchrist, S. & Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American economic review*, 102(4), 1692–1720.

Gilchrist, S. & Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4), 1692–1720.

Gnewuch, M. & Zhang, D. (2025). Monetary policy, firm heterogeneity, and the distribution of investment rates. *Journal of Monetary Economics*, 149, 103721.

Gourio, F. & Kashyap, A. K. (2007). Investment spikes: New facts and a general equilibrium exploration. *Journal of Monetary Economics*, 54, 1–22.

Hennessy, C. A. & Whited, T. M. (2007). How costly is external financing? evidence from a structural estimation. *The Journal of Finance*, 62(4), 1705–1745.

Jeenas, P. (2023). *Firm Balance Sheet Liquidity, Monetary Policy Shocks, and Investment Dynamics*. Working Papers 1409, Barcelona School of Economics.

Jermann, U. & Quadrini, V. (2012). Macroeconomic effects of financial shocks. *American Economic Review*, 102(1), 238–271.

Jiao, F. & Zhang, C. (2022). Lumpy investment and credit risk. *Journal of Corporate Finance*, 77, 102293.

Jordà, Ò., Schularick, M., & Taylor, A. M. (2013). When credit bites back. *Journal of Money, Credit and Banking*, 45(s2), 3–28.

Khan, A. & Thomas, J. K. (2008). Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics. *Econometrica*, 76(2), 395–436.

Khan, A. & Thomas, J. K. (2013). Credit shocks and aggregate fluctuations in an economy with production heterogeneity. *Journal of Political Economy*, 121(6), 1055–1107.

Kiyotaki, N. & Moore, J. (1997). Credit cycles. *Journal of political economy*, 105(2), 211–248.

Kozlowski, J., Veldkamp, L., & Venkateswaran, V. (2019). The tail that keeps the riskless rate low. *NBER Macroeconomics Annual*, 33(1), 253–283.

Kozlowski, J., Veldkamp, L., & Venkateswaran, V. (2020). The tail that wags the economy: Beliefs and persistent stagnation. *Journal of Political Economy*, 128(8), 2839–2879.

Leamer, E. E. (2021). *Why Are Some Recoveries Short and Others Long?* Technical report, National Bureau of Economic Research.

Lee, H. (2022). Striking While the Iron Is Cold: Fragility after a Surge of Lumpy Investments. *Working Paper*, (pp.6̃8).

Lee, H. (2026). A Global Dynamic Nonlinear Solution Framework and the Repeated Transition Method. *Working paper*.

Melcangi, D. (2024). Firms' precautionary savings and employment during a credit crisis. *American Economic Journal: Macroeconomics*, 16(1), 356–386.

Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29(2), 449–470.

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

Ottonello, P. & Winberry, T. (2024). *Capital, Ideas, and the Costs of Financial Frictions*. Working Paper 32056, National Bureau of Economic Research.

Reifschneider, D., Wascher, W., & Wilcox, D. (2015). Aggregate supply in the united states: Recent developments and implications for the conduct of monetary policy. *IMF Economic Review*, 63(1), 71–109.

Rognlie, M., Shleifer, A., & Simsek, A. (2018). Investment hangover and the great recession. *American Economic Journal: Macroeconomics*, 10(2), 113–153.

Senga, T., Thomas, J., & Khan, A. (2017). *Default Risk and Aggregate Fluctuations in an Economy with Production Heterogeneity*. 2017 Meeting Papers 889, Society for Economic Dynamics.

Thomas, J. K. (2002). Is lumpy investment relevant for the business cycle? *Journal of political Economy*, 110(3), 508–534.

Whited, T. M. (2006). External finance constraints and the intertemporal pattern of intermittent investment. *Journal of Financial Economics*, 81(3), 467–502.

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

Xiao, J. (2022). Borrowing to save and investment dynamics. *Available at SSRN 3478294*.

Zwick, E. & Mahon, J. (2017). Tax policy and heterogeneous investment behavior. *Ameri-*

can Economic Review, 107(1), 217–48.