Monetary Policy Without Borrowing: Capacity Constraints and Lumpy Investment*

S. Borağan Aruoba† Andrés Fernández‡ Will Lu§ Felipe Saffie¶

March 2026

Abstract

We study how debt access shapes firm-level monetary transmission using monthly administrative data for the universe of Chilean firms matched to tax records, a comprehensive credit registry, and monetary policy surprises. About 70% of a policy surprise passes through to rates on new loans within 20 days. A 100-basis-point contractionary surprise lowers cumulative investment by about 0.18 percentage points of assets after 24 months, which is roughly 3.2% of capital. This response is concentrated among firms with positive debt at the time of the surprise; firms without debt show near-zero response. Investment also rises by about 0.35 percentage points of assets when firms borrow for the first time and sustain debt. We interpret these facts with a model of limited debt access and lumpy capacity expansion. The model replicates the results in the data where firms that do not have access to debt but have a desire to grow, do not change their plans after a monetary policy change. These firms continue to buy as much capital as they can since their return of doing so far exceeds the user cost due to the lumpy nature of their operations.

Keywords: monetary policy, firm heterogeneity, debt access, financial frictions, lumpy investment.

JEL Codes: D22, D25, E22, E44, E52

*We thank Saleem Bahaj, Martin Ellison, Priit Jeenas, Guillermo Ordoñez, Pablo Ottonello (discussant), Gianluca Violante, and Immo Schott for comments and suggestions. We thank Bernabe Lopez-Martin for an early collaboration. We are grateful to Andrew Castro, Pablo Muñoz, and Eugene Oue for excellent research assistance. The views expressed are those of the authors and do not necessarily represent the views of the International Monetary Fund, its Executive Board, or IMF management, nor those of the Central Bank of Chile or its board members.

†University of Maryland and NBER.

‡International Monetary Fund.

§Central Bank of Chile.

¶University of Virginia (Darden) and NBER.

1 Introduction

Standard models of monetary transmission imply a negative relationship between investment and the interest rate. In the textbook user-cost channel, a higher policy rate raises the cost of capital so firms cut investment. Two things are required for the firms to react to a change in the policy rate: the rate at which the firms are able to borrow needs to react, at least imperfectly, to the change in the policy rate, and the firm should be making a decision at the margin. In this paper we use empirical evidence from Chile and a theoretical model to show that there are circumstances in which firms may not change their investment decisions in response to a change in the policy rate. Specifically, we show that the first condition above is indeed satisfied – firms' borrowing rates quickly adjust after a surprise monetary policy rate change. However, our main empirical result is that firms that do not have debt do not change their investment decisions when faced with a policy rate change.

We study this question using monthly administrative records for the universe of Chilean firms, matched to tax filings, balance-sheet information, a comprehensive credit registry, and monetary policy surprises. The data let us observe both real outcomes and financing positions at high frequency. Four facts summarize the evidence. First, pass-through to the price of credit is rapid: about 70% of a policy surprise is reflected in rates on newly issued loans within 20 days. Second, contractionary surprises reduce investment with delay but with persistent effects: a 100-basis-point surprise lowers cumulative investment by about 0.18 percentage points of assets after 24 months, which amounts to roughly 3.2% of capital. Third, these real effects are concentrated among firms with positive debt, while firms without debt display responses close to zero. Fourth, investment rises sharply when firms borrow for the first time and sustain debt, by about 0.35 percentage points of assets around the borrowing event.

Taken together, these facts suggest that an important margin in monetary transmission is whether firms are borrowing or not. The fourth result is especially informative. Firms that have no debt for an extended period, then borrow for the first time and keep positive debt, display a discrete increase in investment around the borrowing event. This pattern is difficult to reconcile with a view in which debt status is merely a passive correlate of firm characteristics. Instead, it points to a state change: obtaining borrowing capacity relaxes a relevant financing margin for expansion. This helps explain the third result – that firms without debt do not respond. These firms are not making decisions at the margin, rather they are constrained.

To interpret these facts, we develop a heterogeneous-firm model with limited debt access and lumpy capacity expansion. Firms produce with discrete plants, face plant-level capacity

limits, and are subject to span-of-control. Firms without debt access must finance expansion out of retained earnings. They therefore often follow a save-then-jump policy: they operate existing plants at or near capacity, accumulate liquid assets, and expand only when they can afford the next plant and the associated capital. In this region, firms can have marginal products of capital well above the user cost and yet fail to adjust investment on impact after a change in interest rates, because affordability rather than the marginal user cost is binding. Once firms obtain debt access, they choose capital at the margin and can scale immediately to the access-regime target. The model thus links rapid pass-through to loan rates with weak impact responses among firms that are off the relevant financing margin.

This mechanism is distinct from those emphasized in most existing models of firm-level monetary transmission. Ottonello and Winberry (2020) show that firms with low default risk respond more strongly to monetary shocks and interpret this through heterogeneous financing schedules: the relevant heterogeneity is largely among firms that already borrow. Jeenas (2024), Cloyne et al. (2023), Anderson and Cesa-Bianchi (2024), and Bahaj et al. (2022) likewise emphasize liquidity, age, leverage, or collateral as sources of heterogeneous responses within the set of financially active firms. Our evidence points to a more basic margin. Before asking which borrowers are more responsive, it is useful to ask which firms are on the borrowing margin at all. Moreover, our data overcomes some of the restrictions in the existing literature that would prevent one to identify our main result: having quarterly or lower frequency, focusing on publicly-traded firms, which are usually much larger than the generic firm, using a credit registry that has a large minimum for inclusion, making it hard to identify firms without debt.

The model comparison is equally important. Models with smooth borrowing frictions are not naturally designed to deliver our sharp split between firms with debt and firms without it. In those environments, monetary policy changes the marginal cost of funds, and firms remain on a generalized Euler or Kuhn-Tucker condition. Transmission can therefore be heterogeneous, but it is harder to obtain the stark nonresponse we document among firms without debt. Models of constrained accumulation in the spirit of Buera and Shin (2013), Moll (2014), and Midrigan and Xu (2014) are closer, because constrained firms need not respond one-for-one to changes in interest rates. But absent an additional nonconvexity, those environments tend to generate a smoother gradient of responses than the debt/no-debt distinction in our data. A pure fixed-cost Ss model can generate inaction, but there firms effectively finance a deadweight adjustment wedge, and the force is less naturally persistent over the life cycle. Our model combines a genuine extensive borrowing margin with productive and repeated indivisibilities: firms borrow to expand capacity, and self-financing remains slow because growth requires successive discrete increases in scale.

The paper makes three contributions. First, it documents a new fact about monetary transmission using monthly administrative data for the universe of firms in Chile: investment responses to monetary policy are concentrated among firms with positive debt, while firms without debt show little response at short horizons in the baseline specification. Second, it shows that investment rises discretely when firms first obtain debt and sustain it, providing direct evidence that borrowing capacity is tied to lumpy expansion. Third, it provides a tractable model in which debt access determines whether firms are on the relevant financing margin, thereby rationalizing why monetary policy can move borrowing rates quickly while generating weak impact responses for a large set of firms.

Our findings relate to several strands of literature. A large empirical literature studies heterogeneous firm responses to monetary policy, emphasizing default risk, liquidity, collateral, leverage, or age (Ottonello and Winberry, 2020; Jeenas, 2024; Cloyne et al., 2023; Anderson and Cesa-Bianchi, 2024; Bahaj et al., 2022; Durante et al., 2022). Our contribution is to highlight a more basic extensive margin: whether firms are financially active in the first place. The paper also relates to work on lumpy investment and nonconvex adjustment (Cooper and Haltiwanger, 2006; Gourio and Kashyap, 2007), and to recent work linking financing choices to lumpy capacity expansion (Görtz et al., 2023). Relative to that literature, our focus is on how debt access changes the transmission of monetary policy. More broadly, our mechanism is related to models in which financial frictions interact with firm dynamics and self-financing (Buera and Shin, 2013; Moll, 2014; Midrigan and Xu, 2014), and to recent work that studies how lumpy investment shapes the distributional response of firms to monetary policy (Gnewuch and Zhang, 2025). Finally, we build on the literature that uses monetary policy surprises for identification (Gertler and Karadi, 2015; Jarociński and Karadi, 2020; Aruoba et al., 2021) and bring that approach to administrative firm-level data with direct information on both debt and investment.

The rest of the paper is organized as follows. Section 2 describes the data, sample construction, the credit data, and monetary policy surprises. Section 3 presents the empirical evidence on pass-through, investment dynamics, heterogeneity by debt status, and first-time borrowing. Section 4 develops the model. Section 5 provides a quantitative illustration of the mechanism. Appendix A.2 shows that, in the public-firm-with-debt subsample, we recover the leverage gradient emphasized by Ottonello and Winberry (2020).

2 Data, Measurement, and Monetary Policy Surprises

This section describes the monetary policy surprise series, the administrative firm-level data, the sample construction, the credit registry, the loan-origination data, and the construction

of our main variables. The key advantage of the data is that they allow us to observe, at monthly frequency and for the universe of formal firms in Chile, both real outcomes—such as investment and employment—and firms' financing positions, including debt stocks and the interest rates paid on newly issued loans. This combination is central to the paper. It lets us trace monetary transmission from policy surprises to borrowing rates and then to real outcomes, while distinguishing firms with debt from firms without it.

Our main sample spans 2009m1 to 2019m12. The unit of observation is the firm-month. We combine tax records from the Chilean Internal Revenue Service with a comprehensive credit registry compiled from supervisory and central bank sources. After sample restrictions, the resulting panel contains 4,718,460 firm-month observations for 57,459 firms.

2.1 Monetary policy surprises

We identify exogenous movements in monetary policy using the surprise series in Aruoba et al. (2021). For each monetary policy meeting of the Central Bank of Chile, the surprise is defined as the realized policy decision minus the market expectation recorded immediately before the meeting. Formally, for a meeting at time t,

$$MPS_t \;=\; MPR_t - MPR_t^e, \tag{1}$$

where MPR_t is the announced monetary policy rate and MPR_t^e is the median Bloomberg expectation prior to the announcement.

This measure has two advantages for our setting. First, it isolates the unanticipated component of the policy decision rather than movements that were already priced in. Second, because Bloomberg expectations can be updated close to the announcement, the series mitigates concerns that arise with lower-frequency expectation measures. Over our sample period, we observe 15 non-zero surprises: 5 positive and 10 negative, with a modal size of 25 basis points. Figure 1 plots the series.

Throughout the paper, we denote the surprise by $\varepsilon_t \equiv MPS_t$, measured in percentage points. Thus, $\varepsilon_t = 0.25$ corresponds to a 25-basis-point surprise tightening.

2.2 Administrative firm data and credit registry

We combine four administrative sources.

First, we use monthly VAT filings (Form F29) from the Chilean Internal Revenue Service (SII). These records provide firm-level information on sales and purchases, including purchases of machinery and equipment, which we use to measure investment.

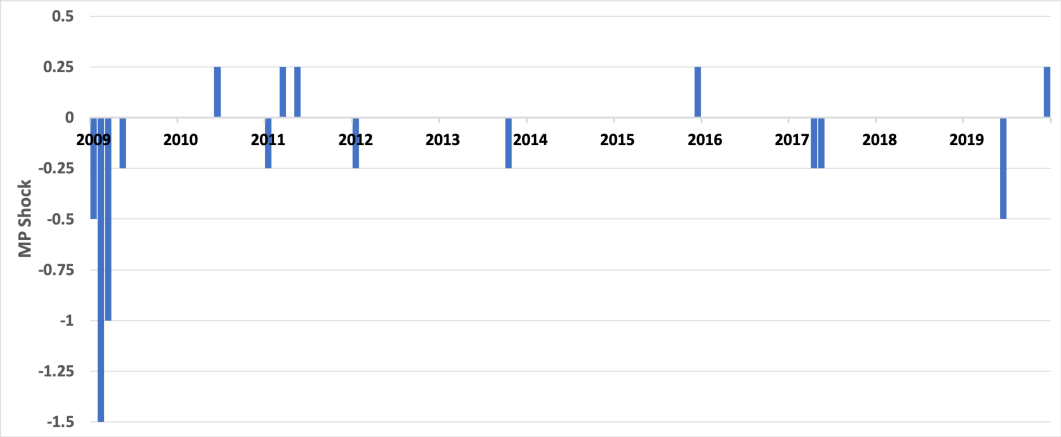


Figure 1: Monetary policy surprises constructed from Bloomberg expectations.

Second, we use matched employer-employee records (Form DJ1887), also from the SII. These data provide employment and wage information at monthly frequency.

Third, we use annual corporate income tax filings (Form F22), which contain balance-sheet information, including total assets. We use these filings to construct scale variables and financial ratios.

Fourth, we merge these tax data to a comprehensive credit registry built from information collected by the Financial Markets Commission (CMF) and the Central Bank of Chile. The registry contains monthly debt stocks and loan-level originations. The debt-stock files record outstanding bank debt and distinguish normal from delinquent amounts. The loan-origination files report the amount, maturity, and contractual interest rate on newly issued bank loans. We also observe monthly outstanding bond liabilities, which we use when constructing total debt measures.

The merge is performed using a common anonymized tax identifier. The result is a monthly panel that combines firms' real outcomes, balance-sheet characteristics, debt positions, and borrowing terms.

2.3 Main variables and debt status

Our main real outcome is investment in machinery and equipment, measured from monthly VAT filings. In the empirical analysis, we scale investment by total assets from the annual tax filing. We also use employment, sales, age, leverage, and delinquent debt shares as controls in the local projection specifications.

The central financial variable in the paper is debt status. We define

$$D_{jt} \;=\; \mathbb{1}\{\text{firm } j \text{ has positive outstanding bank debt in month } t\}. \tag{2}$$

This is intentionally a simple measure. Our main empirical distinction is not between firms with different intensities of leverage, but between firms that are financially active in the debt market and firms that are not. We interpret positive outstanding debt as the empirical proxy for being on or near the borrowing margin. In the model, this empirical distinction maps into the access versus no-access state. We also construct total debt, combining bank debt and bonds, leverage as total debt over assets, and a delinquent debt ratio.

Several descriptive facts motivate this focus. In a typical month, roughly 75% of firms in our sample have positive debt. At the same time, debt status is far from universal or fixed. About 8.5% of firms never have debt during the sample, while about 40% always have debt. Among firms born in 2010, roughly 60% enter without debt, 40% enter with debt, and about 76% have obtained debt within their first five years. These patterns indicate that debt status is both economically important and dynamic over the firm life cycle.

2.4 Sample construction

We impose a set of restrictions to focus on regular operating firms with sufficiently reliable reporting.

First, we exclude utilities and public administration. Second, we require firms to have at least three employees and positive sales. Third, we require non-missing asset information in the annual tax filing. Fourth, we keep firms that report in at least 75% of the months they are potentially present in the sample, with at least one run of 12 consecutive reporting months. These restrictions remove irregular reporters and firms with sparse administrative histories. Finally, we winsorize key variables and ratios at the top and bottom 0.5%.

After these restrictions, the final panel contains 4,718,460 firm-month observations for 57,459 firms between 2009m1 and 2019m12. The median firm has 14 employees. Table 1 reports summary statistics.

2.5 Loan origination data

To study pass-through from monetary policy surprises to firms' borrowing costs, we use the loan-origination files from the credit registry. These data cover newly issued bank loans from 2012m4 to 2019m12 and report the contractual interest rate, maturity, amount, and loan type. This loan-level information is essential for Result 1, because it lets us measure the price of credit faced by firms directly rather than inferring it from balance-sheet outcomes.

The flow dataset is large and comprehensive. It contains 7,106,021 loan-level observations over the sample period. As with the firm-level panel, we winsorize key loan variables at the top and bottom 0.5%.

Table 1: Summary statistics

	Mean	Median	S.D.	95th Perc.
Sales*	324	43	1308	1142
Capital stock*	941	27	26000	1224
Total assets*	7666	434	41057	22034
Investment*	4.8	0	30	14
Share with positive debt (%)	75.0	–	–	–
Age	12	10	7.1	24
No. employees	58	14	165	230
Normal debt stock*	509	16	2351	1838
Bad debt stock*	2.8	0	24	0
Total debt*	512	17	2352	1844
Sales-to-assets ratio	0.2	0.1	0.3	0.6
Leverage ratio	0.2	0.1	0.3	0.7
No. firms	57,459			
Observations	4,718,460			

Notes: Summary statistics for 2009m1–2019m12. *Currency-denominated variables are in millions of Chilean pesos (CLP). Positive investment occurs in 20.6% of firm-month observations. The statistics on never, always, and first-time debt use firm histories in the final sample.

The next section uses these data to establish four facts: rapid pass-through from policy surprises to loan rates, delayed but persistent investment responses, heterogeneity in those responses by debt status, and a sharp increase in investment when firms borrow for the first time.

3 Firm-Level Evidence on Monetary Transmission

This section documents four facts about firm-level monetary transmission. We first show that monetary policy surprises pass through rapidly to the interest rates firms pay on new loans. We then show that investment responds with delay but with persistent effects. The central empirical result is that these real responses are concentrated among firms with positive debt, while firms without debt display little response at short horizons in our baseline specification. Finally, we show that investment rises sharply when firms borrow for the first time and sustain debt. Taken together, these facts motivate a distinction between firms that are already adjusting on the borrowing margin and firms that are not.

Throughout, we denote the monetary policy surprise by ε_t, with a positive value corresponding to a contractionary surprise.

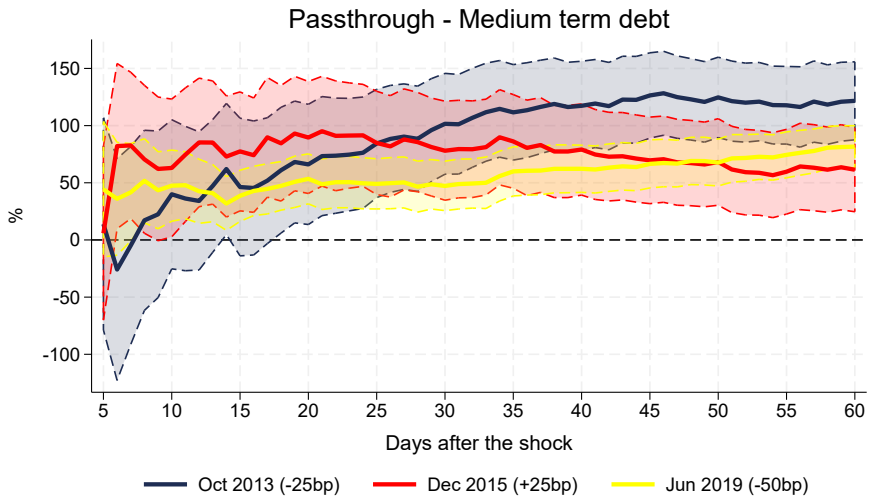


Figure 2: **Pass-through to loan rates.** Event-study estimates of the response of contractual loan rates to selected monetary policy surprises, scaled by surprise size.

3.1 Result 1: Pass-through from monetary policy surprises to loan rates

We begin with the price of credit. If monetary policy does not pass through to firms' borrowing costs, then there is little reason to expect firm-level investment to respond through a financial channel. To study pass-through, we use the loan-origination data described in Section 2 and focus on three surprise episodes with clean post-announcement windows in the loan-flow data: October 2013, December 2015, and June 2019.

For each surprise date s, let $r_{j\tau}$ denote the annualized contractual interest rate on a newly issued loan to firm j on day τ. For each horizon $d \in \{5, \ldots, 60\}$, we estimate

$$\ln(1+r_{j\tau}) = \alpha_j + \beta_d \, \mathbb{1}\{\tau > s\} + \Gamma X_{j\tau} + u_{j\tau}, \qquad \tau \in [s-30, \ldots, s-1, s+1, \ldots, s+d], \quad (3)$$

where α_j are firm fixed effects and $X_{j\tau}$ includes loan controls such as maturity and amount. Standard errors are clustered at the firm level. In practice, the figures report the post-estimation scaled coefficient $\hat{\beta}_d/\varepsilon_s$, so the estimates can be interpreted as pass-through per percentage-point surprise.

Figure 2 shows rapid pass-through. Within roughly 20 days, about 70% of the surprise is reflected in the rates firms pay on newly issued loans. The first conclusion of the paper is therefore straightforward: monetary policy shocks are transmitted quickly to the price of external finance faced by borrowing firms.

3.2 Result 2: Baseline investment response

We next ask whether monetary policy surprises affect firms' real investment. Our outcome is cumulative investment in machinery and equipment, scaled by total assets. For horizon $h \geq 0$, define

$$Y_{j,t}^{I}(h) \equiv \sum_{m=0}^{h} \left(\frac{i_{j,t+m}}{A_{j,t+m}} \right), \tag{4}$$

where i_{jt} is machinery and equipment investment and A_{jt} is total assets. Thus, the horizon-0 outcome is the investment-to-assets ratio in the shock month. Since the shocks are dated by meeting month and the outcome is monthly, the medium-run responses are the relevant objects for interpretation.

We estimate local projections of the form

$$Y_{j,t}^{I}(h) = \alpha_j + \alpha_{s,q(t)} + \beta_h \varepsilon_t + \Gamma X_{jt} + u_{j,t}(h), \tag{5}$$

where α_j are firm fixed effects, $\alpha_{s,q(t)}$ are industry-by-quarter fixed effects, and X_{jt} includes firm age, employment, the sales-to-assets ratio, leverage, and the delinquent-debt ratio. Standard errors are two-way clustered at the firm level and the industry-quarter level.

Figure 3 shows a delayed but persistent response. A contractionary surprise lowers cumulative investment, with the effect becoming statistically clear around 11 months and stabilizing around 23 months. At the 24-month horizon, a 100-basis-point contractionary surprise lowers cumulative investment by about 0.18 percentage points of assets. Using the approximation in Appendix D, this corresponds to roughly a 3.2% decline in the machinery-and-equipment capital stock after two years.

Two points are worth emphasizing. First, the response is not immediate, even though pass-through to borrowing rates is fast. Second, the effect is persistent once it appears. This combination already suggests that monetary transmission at the firm level is not well summarized by a frictionless static user-cost calculation.

3.3 Result 3: Debt status and heterogeneous investment responses

We now turn to the main empirical question of the paper: are investment responses concentrated among firms with debt? To answer this, we interact the monetary policy surprise with our measure of debt status:

$$Y_{j,t}^{I}(h) = \alpha_j + \alpha_{s,q(t)} + \beta_h \varepsilon_t + \beta_{h,D} \left(\varepsilon_t \times D_{jt} \right) + \Gamma X_{jt} + u_{j,t}(h), \tag{6}$$

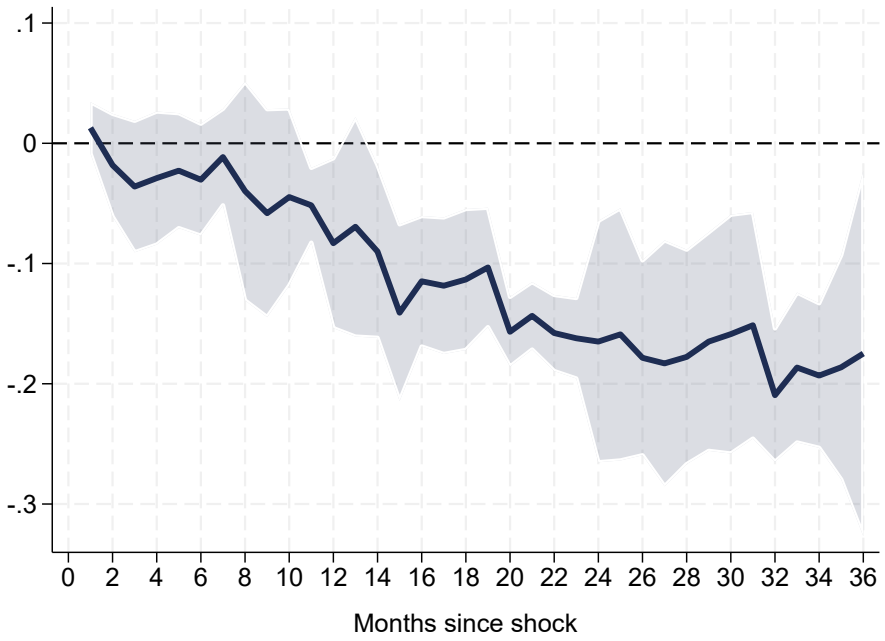


Figure 3: **Baseline cumulative investment response.** Local projection estimates of the effect of a monetary policy surprise on cumulative investment.

where $D_{jt} = 1$ if firm j has positive outstanding bank debt in month t.

Figure 4 shows a sharp difference by debt status. The investment response is driven by firms with debt. Firms without debt display responses close to zero at short horizons in the baseline specification, while firms with debt show a substantially larger and more persistent decline in investment after a contractionary surprise.

This result is central. Much of the existing literature studies heterogeneous responses among firms that are already financially active, emphasizing leverage, collateral, or balance-sheet strength. Our evidence points to a more basic extensive margin. Before asking which borrowers respond more, one needs to ask which firms are on the borrowing margin at all. In our data, that distinction is quantitatively important. Appendix A.2 shows that leverage gradients remain present in our data, and that when we restrict attention to public firms with debt we recover the qualitative pattern in Ottonello and Winberry (2020).

Because debt status is measured in the month of the policy meeting, equation (6) is best read as classifying firms by financing regime at the time of the shock. The first-borrowing event study below is useful precisely because it separately isolates discrete transitions into debt. At longer horizons, the no-debt estimates are less precise and need not remain exactly zero. In the model, that is natural: changes in saving behavior can affect when no-access firms reach the next lumpy expansion threshold, so the response of firms without borrowing capacity need not be literally zero forever even if the impact response is weak.

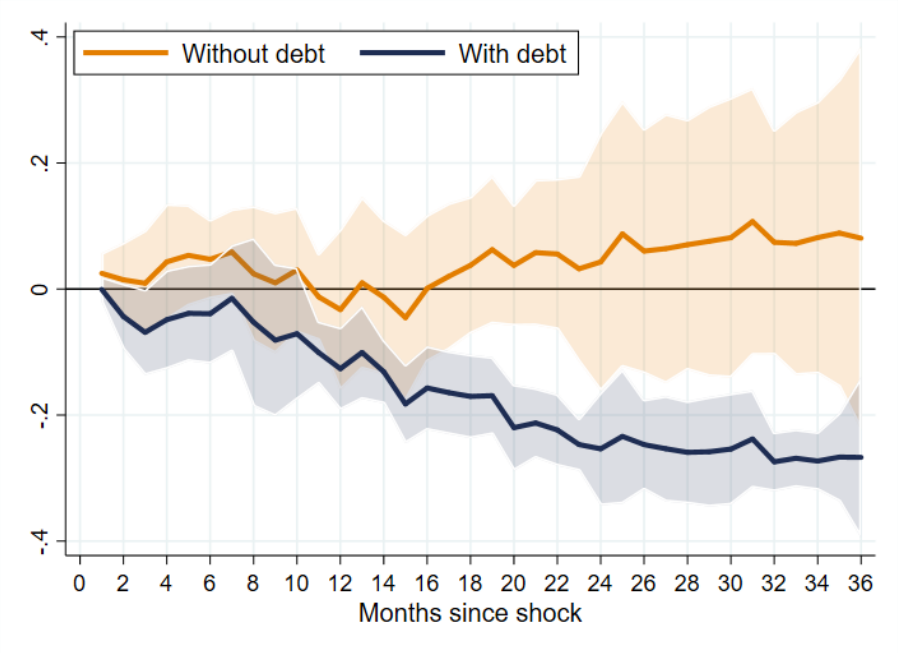


Figure 4: **Debt status and monetary transmission.** Local projection responses for firms with and without outstanding bank debt.

3.4 Result 4: Investment around first-time borrowing

The debt-interaction results show that firms with and without debt respond differently to monetary policy. We now ask whether investment changes discretely when firms move from one state to the other. To do so, we study firms that borrow for the first time.

We define an event at time $t = 0$ when a firm satisfies four conditions: (i) it has no debt in the previous 18 months; (ii) it is observed for at least 12 months before and 24 months after the event; (iii) it takes on debt at $t = 0$; and (iv) it maintains positive debt for the following four months. The event-study sample contains 2,328 firms and 254,033 firm-month observations.

We estimate

$$\frac{i_{jt}}{A_{jt}} = \alpha_j + \alpha_t + \sum_{k=-12}^{24} \beta_k D_{kjt} + \Gamma X_{jt} + e_{jt}, \tag{7}$$

where α_j are firm fixed effects, α_t are time fixed effects, D_{kjt} is an indicator for event time k, and X_{jt} includes age, employment, and the sales-to-assets ratio. Standard errors are two-way clustered by firm and calendar time. Coefficients are normalized relative to event time $k = -1$.

Figure 5 shows a sharp increase in investment around the borrowing event. Investment rises discretely when firms obtain debt and keep it, by roughly 0.35 percentage points of

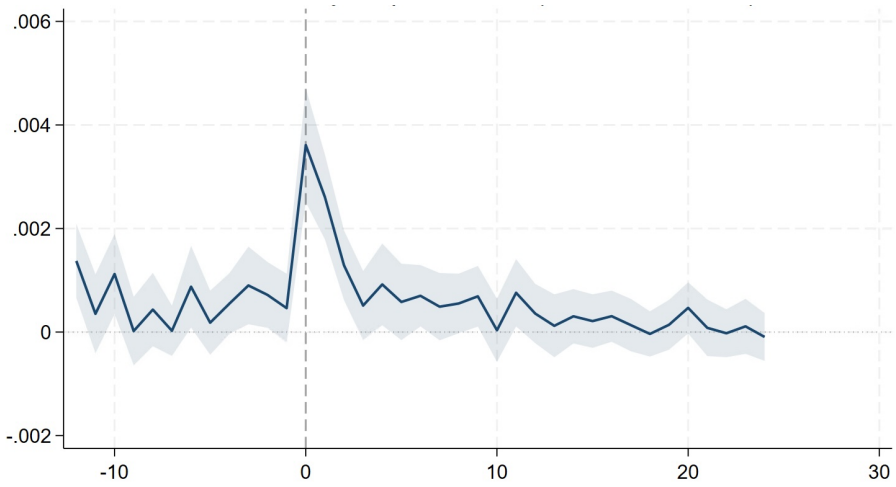


Figure 5: **Investment around first-time borrowing.** Event-study estimates of investment relative to the month of first sustained borrowing.

assets relative to the pre-event period. Pre-event coefficients are economically small and do not display a pronounced trend. This pattern is consistent with a lumpy expansion margin: firms without debt appear to accumulate resources and then expand when borrowing becomes feasible. The estimation window is $[-12, 24]$; the figure displays the central portion of that window for readability.

This event-study evidence is important for interpretation. The debt-status interaction in Figure 4 could in principle reflect persistent differences in firm types. Figure 5 instead points to a state change. When firms move from no debt to positive debt, investment rises sharply. That is precisely the pattern one would expect if debt status proxies for a relevant financing constraint for discrete expansion.

Taken together, these four facts motivate the model that follows. Monetary policy passes through rapidly to firms' borrowing costs, but real responses are concentrated among firms with debt, while firms without debt show little response at short horizons and invest sharply when they first obtain debt. The next section develops a model in which these patterns arise because debt access determines whether firms are adjusting at the margin or saving toward the next lumpy expansion.

4 A Model of Debt Access and Lumpy Capacity Expansion

This section develops a simple model that rationalizes the four empirical facts in Section 3. The key ingredients are: (i) firms expand through discrete plants, (ii) each plant has a

capacity limit, (iii) span-of-control makes the profitability of capital depend on firm scale, and (iv) debt access is limited. The model implies that some firms adjust capital at the margin and satisfy the standard user-cost condition, while others are off that margin and instead save toward the next lumpy expansion.

We keep the environment deliberately parsimonious. The objective is not to build a rich theory of financial contracting, but to isolate a simple mechanism through which debt access changes the firm-level transmission of monetary policy.

4.1 Environment and timing

Time is discrete. A firm is characterized by permanent productivity z, liquid asset holdings b, capital per plant k, the number of plants n, and a debt-access status $a \in \{N, A\}$, where N denotes no access and A denotes access. All firms enter without debt access.

A firm with productivity z operates $n \in \mathbb{N}$ plants. Output is produced with

$$Y(z, n, k, \ell) = z\, n^{\xi} k^{\alpha_k} \ell^{\alpha_\ell}, \qquad \alpha_k, \alpha_\ell \in (0,1), \quad \alpha_k + \alpha_\ell < 1, \quad \xi \leq 1, \tag{8}$$

where k and ℓ denote capital and labor *per plant*. The parameter $\xi \leq 1$ captures span-of-control: adding plants raises output, but less than proportionally when $\xi < 1$. Each plant faces a capacity limit

$$k \leq \bar{k}. \tag{9}$$

Plants are discrete and irreversible. If the firm starts the period with n plants, next period it can choose

$$n' \in \{n, n+1, \ldots\}, \tag{10}$$

and adding a plant requires paying a fixed cost $F > 0$ per plant.

Firms can save or borrow using a one-period risk-free bond. Let b' denote next period's bond position, with $b' > 0$ corresponding to savings and $b' < 0$ to debt. The gross real interest rate is $1 + R$, capital depreciates at rate $\delta \in (0,1)$, and $\beta \in (0,1)$ denotes the discount factor. We normalize discounting so that

$$\beta(1 + R) = 1. \tag{11}$$

Debt access is the key financial friction. Firms with access can choose $b' \in \mathbb{R}$, subject only to the standard no-Ponzi condition

$$\liminf_{T \to \infty} \beta^T b_{t+T} \geq 0 \qquad \text{along any feasible plan.} \tag{12}$$

Firms without access cannot borrow and must satisfy

$$b' \geq 0. \tag{13}$$

To pin down payout policy in the numerical implementation, we also impose a large upper bound $\bar{b}$ on savings. This upper bound never binds for the real choices (n', k') in the experiments reported below; it only eliminates irrelevant indeterminacy between retained liquid wealth and dividends under (11).

The timing within a period is:

(i) The firm enters the period with state (a, b, k, n, z).

(ii) If $a = N$, the firm knows that, conditional on survival, it will wake up with access next period with probability $p(W)$, where current book value is

$$W \equiv b + nk + Fn. \tag{14}$$

Since W is predetermined, current choices do not affect whether the firm wakes up with access tomorrow.

(iii) The firm produces, chooses labor, and obtains cash on hand.

(iv) The firm chooses (n', k', b', d).

(v) Exit occurs with probability $1 - \zeta$ after choices and before next production. An exiting firm liquidates its chosen next-period assets at

$$\Phi(b', k', n') \equiv b' + n'k'. \tag{15}$$

The access probability is

$$p(W) = \bar{p}\left(1 - e^{-\kappa W}\right), \qquad \bar{p} \in (0, 1], \quad \kappa > 0. \tag{16}$$

Once access is obtained, it is permanent.

The object W should be interpreted as a reduced-form screening statistic rather than as liquidation value. A lender is more willing to extend credit when the entrepreneur has more "skin in the game" in the business, whether that appears as retained liquid wealth or as installed business scale. Including Fn in W captures that idea: a dollar committed to building organizational scale counts toward future access, even if it is not separately

recoverable at liquidation. Using $W = b + nk + Fn$ also avoids an artificial incentive to hold liquid wealth rather than productive business capital purely to raise the access probability.

4.2 Static profits, MPK, and target plant size

Given (z, n, k) and wage w, the firm chooses labor within the period:

$$\pi(z, n, k; w) \equiv \max_{\ell \geq 0} \left\{ z\, n^{\xi} k^{\alpha_k} \ell^{\alpha_\ell} - wn\ell \right\}. \tag{17}$$

The next proposition gives the closed-form profit function after labor is chosen optimally.

Proposition 4.1 (Static labor demand and profits)**.** *Given (z, n, k) and wage w, optimal labor per plant is*

$$\ell(z, n, k; w) = \left(\frac{\alpha_\ell z}{w} n^{\xi - 1} k^{\alpha_k} \right)^{\frac{1}{1-\alpha_\ell}}, \tag{18}$$

and operating profits are

$$\pi(z, n, k; w) = (1 - \alpha_\ell)\, z^{\frac{1}{1-\alpha_\ell}} \left(\frac{\alpha_\ell}{w} \right)^{\frac{\alpha_\ell}{1-\alpha_\ell}} n^{\frac{\xi - \alpha_\ell}{1-\alpha_\ell}} k^{\frac{\alpha_k}{1-\alpha_\ell}}. \tag{19}$$

Define total capital as $K \equiv nk$. We define the marginal product of total capital by

$$MPK(z, n, k; w) \equiv \frac{\partial \pi(z, n, k; w)}{\partial K} = \frac{1}{n} \frac{\partial \pi(z, n, k; w)}{\partial k}. \tag{20}$$

Lemma 4.2 (Properties of the marginal product of capital)**.** *For fixed (z, w), the function $MPK(z, n, k; w)$ is continuous and strictly decreasing in k on $(0, \bar{k}]$. If $\xi < 1$, it is strictly decreasing in n as well. Consequently, whenever the equation*

$$MPK(z, n, k; w) = R + \delta \tag{21}$$

admits a solution, that solution is unique.

These properties imply a well-defined interior target capital per plant. Let $\hat{k}(z, n, R; w)$ solve

$$MPK(z, n, \hat{k}; w) = R + \delta. \tag{22}$$

Then the target capital per plant is

$$k^*(z, n, R; w) = \min\{\hat{k}(z, n, R; w), \bar{k}\}. \tag{23}$$

Figure 6 illustrates the technology side of the mechanism. For firms with more plants,

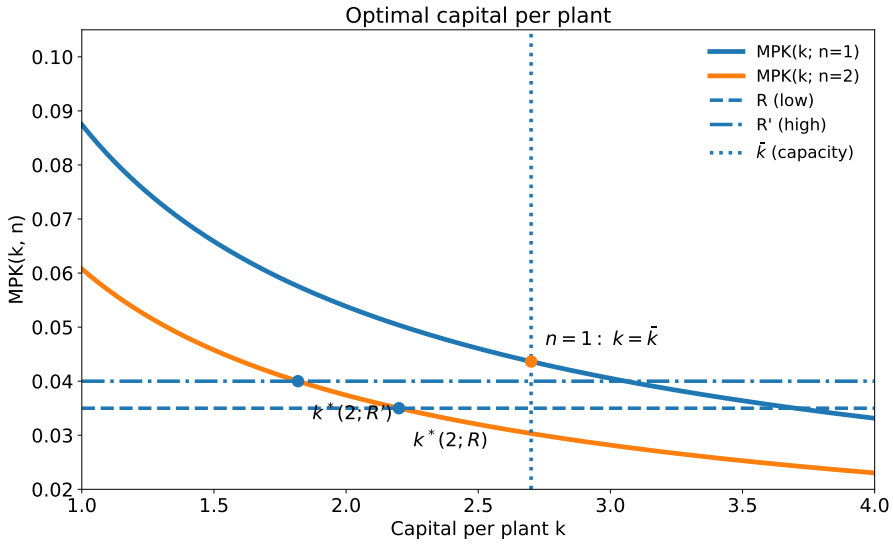


Figure 6: **MPK, target capital, and plant capacity.** Span-of-control and plant-level capacity create a region in which firms are off the standard interior condition $MPK = R + \delta$. A higher interest rate shifts the interior target left, but a constrained firm at $k = \bar{k}$ does not adjust capital on impact.

the MPK schedule lies lower because of span-of-control. If the firm is interior, a rise in the interest rate from R to R' lowers the target capital per plant from $\hat{k}(z, n, R; w)$ to $\hat{k}(z, n, R'; w)$. But if the firm is already at the plant-level capacity limit, the increase in the interest rate does not change capital on impact: the firm remains at $k = \bar{k}$ and therefore off the standard interior Euler condition.

Figure 6 is not yet the full mechanism of the paper. By itself, it shows only that the technology creates a meaningful corner region. Debt access matters because it determines whether firms can escape that region by scaling discretely.

4.3 Budget constraint and dynamic problem

Let X denote cash on hand after production:

$$X = (1 + R)b + \pi(z, n, k; w) + (1 - \delta)nk. \tag{24}$$

The firm's budget constraint is

$$d + b' = X - n'k' - F(n' - n), \qquad d \geq 0, \tag{25}$$

where d denotes dividends. The firm chooses (n', k', b') subject to (9), (10), and the relevant financing constraint.

Let $V^A(z,b,k,n)$ denote the value of a firm with debt access and $V^N(z,b,k,n)$ the value of a firm without access. Then the access firm's problem is

$$V^A(z,b,k,n) = \max_{\substack{n',k',b',d \\ n'\geq n,\ 0\leq k'\leq \bar{k} \\ b'\in\mathbb{R},\ d\geq 0}} \left\{ d + \beta \left[\zeta V^A(z,b',k',n') + (1-\zeta)\Phi(b',k',n') \right] \right\}, \tag{26}$$

subject to (25) and (12).

The no-access firm's problem is

$$V^N(z,b,k,n) = \max_{\substack{n',k',b',d \\ n'\geq n,\ 0\leq k'\leq \bar{k} \\ b'\geq 0,\ d\geq 0}} \left\{ d + \beta \left[\zeta \widetilde{V}(z,b',k',n';W) + (1-\zeta)\Phi(b',k',n') \right] \right\}, \tag{27}$$

subject to (25), where

$$\widetilde{V}(z,b',k',n';W) = p(W)V^A(z,b',k',n') + (1-p(W))\,V^N(z,b',k',n'). \tag{28}$$

Because W is predetermined, the continuation weights in (28) are fixed at the time choices are made.

4.4 Firms with debt access

The role of debt access is easiest to see by starting from the benchmark regime. A firm with access can choose its bond position freely. It can therefore finance a discrete increase in the number of plants immediately rather than waiting to accumulate internal funds.

Proposition 4.3 (Access regime: real-choice independence and immediate access-regime adjustment)**.** *For an access firm:*

(i) the ranking of feasible real choices (n',k') is independent of current liquid wealth b and current installed capital k;

(ii) from any access state, the firm chooses the access-regime optimal real choice immediately rather than waiting to self-finance it;

(iii) if the optimal choice satisfies $k' \in (0,\bar{k})$, then

$$MPK(z,n',k';w) = R + \delta. \tag{29}$$

The proposition says exactly what we need for the paper. Once the firm has debt access, current internal funds no longer constrain next-period real choices. Access restores marginal

adjustment. In particular, if the access-regime optimum lies above the current scale, the firm jumps to it immediately rather than waiting to self-finance.

4.5 Firms without debt access

A firm without debt access faces a different problem. It may wish to operate at a larger scale, but it cannot finance the next plant or the desired capital per plant by borrowing. As a result, a finite number of regimes emerge.

The first is an *interior saving regime*. If the firm can afford its target choice and still carry positive savings forward, then it behaves much like an unconstrained firm and chooses capital to satisfy the usual condition $MPK = R + \delta$.

The second is a *cash-poor regime*. If the firm cannot afford the desired next-period scale and capital, then it sets $b' = 0$ and allocates all available resources to current scale and capital. Investment is then pinned down by affordability rather than by the user cost.

The third is a *capacity regime*. A firm may choose $k' = \bar{k}$ for a given plant choice. In that case, the plant-level capacity constraint binds and any immediate response must come through n' rather than k'.

Proposition 4.4 (No-access regimes)**.** *Fix a no-access firm with state (z, b, k, n) and let (n', k', b', d) be an optimal choice.*

(i) If $b' > 0$ and $k' \in (0, \bar{k})$, then

$$MPK(z, n', k'; w) = R + \delta. \tag{30}$$

(ii) If $b' = 0$, then there exists an optimal representative with $d = 0$, and for such an optimum the budget constraint binds as

$$n'k' + F(n' - n) = X. \tag{31}$$

(iii) If $b' = 0$, $k' \in (0, \bar{k})$, and the interior target for the chosen n' is not affordable, then

$$MPK(z, n', k'; w) > R + \delta. \tag{32}$$

(iv) If $k' = \bar{k}$, the plant-level capacity constraint binds.

This proposition is the core of the model. The no-access firm can be off the standard Euler condition not because the user cost is irrelevant, but because the desired choice is

infeasible. Such a firm would like to operate at a larger scale or with more capital per plant, but the inability to borrow forces it to save toward that target over time. In that sense, the relevant margin is not marginal adjustment but threshold crossing.

4.6 Impact effects of monetary policy

The empirical results show rapid pass-through to borrowing rates but weak short-run responses for firms without debt. The model delivers exactly this pattern.

Corollary 4.5 (Impact response to an interest-rate increase)**.** *Consider an unanticipated increase in the interest rate from R to $R' > R$, holding the wage fixed.*

(i) If a firm chooses interior per-plant capital with slack savings, then the new choice satisfies

$$MPK(z, n', k'; w) = R' + \delta. \tag{33}$$

(ii) If a no-access firm is in the cash-poor regime, then its current-period choice is governed by affordability rather than by the marginal user cost. In particular, if the shock does not change the locally optimal discrete plant choice, current investment need not adjust on impact.

(iii) If a firm is at the capacity limit $k' = \bar{k}$, per-plant capital does not adjust on impact. Any immediate response must come through the discrete plant margin.

The intuition is simple. Access firms, and no-access firms that are already saving on the interior, adjust through the standard user-cost channel. By contrast, no-access firms in the cash-poor region are not choosing capital at the margin. A higher interest rate changes the value of waiting, but if the current-period choice is constrained by the inability to finance the next lumpy expansion, observed investment need not move immediately.

4.7 Mapping the model to the empirical facts

The model organizes the empirical evidence in Section 3 around a single distinction: whether the firm is on the relevant financing margin.

First, Result 1 shows that monetary policy surprises pass through rapidly to the interest rates firms pay on newly issued loans. In the model, this corresponds to a movement in R that immediately affects the user cost of capital for firms that can borrow.

Second, Result 2 shows that aggregate firm investment responds with delay but with persistent effects. In the model, this occurs because the interest-rate change affects some

firms immediately and others only as they move across access states or accumulate sufficient internal funds to cross the next expansion threshold.

Third, Result 3 shows that the investment response is concentrated among firms with debt. The model generates this directly: firms with debt choose capital at the margin and respond to the interest rate, while many firms without debt are in the cash-poor regime and therefore display little impact response.

Fourth, Result 4 shows that investment rises sharply when firms borrow for the first time and sustain debt. In the model, that event corresponds to a transition from the save-then-jump region to one in which the firm can finance the next plant immediately. Because expansion is productive and discrete, investment jumps when access arrives.

4.8 Why this model? A comparison with nearby alternatives

The question is not whether other models can generate *some* heterogeneity in investment responses to monetary policy. Many can. The question is whether they naturally generate the specific pattern we document: rapid pass-through to borrowing rates, little short-run response among firms without debt, and a sharp increase in investment when firms first obtain debt.

A natural benchmark is a model with smooth borrowing frictions. In a broad class of such environments, firms face a state-dependent external finance wedge and the relevant marginal condition can be written schematically as

$$MPK = R + \delta + \tau_i, \tag{34}$$

where τ_i depends on leverage, default risk, collateral, or the shape of the firm's borrowing schedule. A monetary tightening changes R and may also change τ_i, so these models deliver heterogeneous responses across financially active firms. But firms still adjust on a generalized marginal condition. This makes them well suited to studying heterogeneity among borrowers, as in Ottonello and Winberry (2020), but less suited to generating the stark near-zero impact responses we find among firms with no debt.

Models of constrained accumulation in the spirit of Buera and Shin (2013), Moll (2014), and Midrigan and Xu (2014) are closer. Constrained firms may be unable to equate marginal products across uses, and self-financing matters for growth. But absent an additional nonconvexity, these models tend to produce a relatively smooth gradient of responses. Firms gradually relax constraints through retained earnings, but there is no sharp distinction between being able or unable to finance the next discrete expansion.

A second nearby benchmark is a pure fixed-cost Ss model with no borrowing. Let a firm

start with capital k_0 and internal funds w. It can either keep k_0, or pay a fixed cost F and choose a new capital stock k':

$$V = \max \left\{ \pi(k_0) + Rw, \max_{k' \leq w - F} [\pi(k') + R(w - F - k')] \right\}. \tag{35}$$

Conditional on adjustment, the optimal choice solves

$$\pi'(k') = R, \tag{36}$$

and the firm adjusts only if

$$\max_{k' \leq w - F} [\pi(k') - Rk'] - F \geq \pi(k_0). \tag{37}$$

This setup can certainly generate inaction. Small changes in R need not trigger adjustment if the firm remains inside the inaction band. But the object being financed is a deadweight adjustment cost. Once the firm accumulates enough resources to pay F and move near target, the force weakens. It is therefore harder to generate the persistent constrained life-cycle region that emerges in our data and in our model.

Our model adds two ingredients. First, the extensive financial margin is genuine: firms differ in whether they can borrow at all. Second, the nonconvexity is productive and repeated: firms borrow not to pay a wedge, but to add capacity. Because growth requires successive discrete plant expansions, self-financing remains slow and some firms stay in the save-then-jump region for long periods. That combination is what makes the model well suited to the empirical facts in Section 3.

The next section quantifies these forces. We calibrate the model, illustrate firms' life-cycle dynamics with and without debt access, and study a monetary policy exercise in which the aggregate response depends on the distribution of firms across access and adjustment regimes.

Full proofs are provided in Appendix B.

5 Quantitative Illustration

This section illustrates the mechanism of the model and shows how it maps into the cross-sectional responses highlighted in Section 3. We proceed in three steps. We first present a simple calibration and describe the stationary distribution. We then use simulated life-cycle paths to show how firms with and without debt access evolve differently over time. Finally, we study a monetary policy exercise in which the real interest rate rises temporarily and

Table 2: Baseline calibration

Object	Symbol(s)	Value
Technology / span of control	$(\alpha_k, \alpha_\ell, \xi)$	(0.60, 0.30, 0.95)
Depreciation	δ	0.025
Real interest rate (baseline)	R	0.025
Labor supply	$\bar{L}$	0.57
Cost of building a plant	F	0.24
Exit probability	$1-\zeta$	0.01
Per-plant capacity	$\bar{k}$	1.2
Access probability	$p(W)$	$(\bar{p}, \kappa) = (0.5,\ 0.05)$
Productivity distribution	z	$\{0.67, 0.695, 0.72, 0.745, 0.77\}$
Probability mass		$(0.14, 0.23, 0.28, 0.21, 0.15)$

examine both aggregate responses and heterogeneity in impact effects across firms.

The purpose of this section is illustrative rather than fully quantitative. We do not claim to match the full set of moments in the Chilean data. Instead, we show that the model generates the qualitative pattern central to the paper: monetary policy passes through immediately to the relevant financing margin, but impact responses are concentrated among firms with debt access or among firms already adjusting on the interior margin.

5.1 Calibration and stationary equilibrium

Table 2 reports the baseline parameterization. The span-of-control and production parameters determine how the marginal product of capital varies with firm scale. The plant cost F and the capacity limit $\bar{k}$ determine the size of nonconvexities. The access-arrival function governs how likely firms without debt are to obtain debt access as they accumulate scale and resources.

We solve for a stationary equilibrium in which the wage clears the labor market,

$$\bar{L} = \int n\,\ell(z, n, k; w)\,d\mu(z, b, k, n, a), \tag{38}$$

where μ is the invariant distribution over firm states. The productivity distribution is discrete and chosen to generate heterogeneity in desired firm scale. In the stationary distribution generated by this calibration, 78% of firms have debt access, 82% are at target scale, and 18% of no-access firms remain below target scale. That last group is the quantitative counterpart of the save-then-jump region highlighted in the empirical analysis.

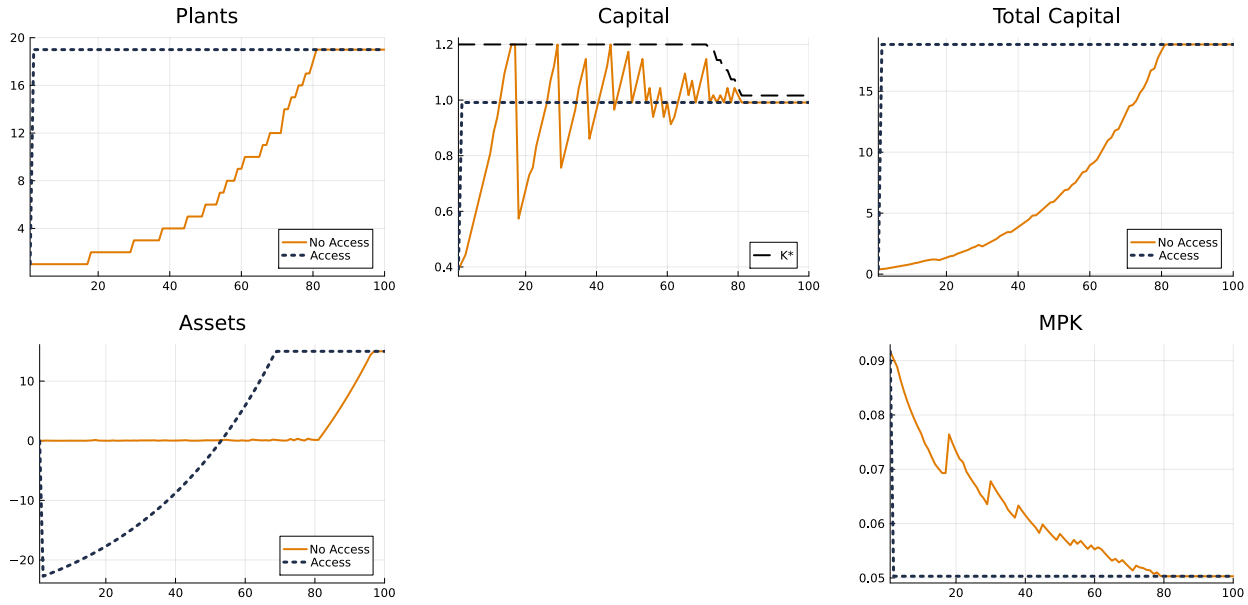


Figure 7: **Life cycle of a high-productivity firm: access versus no access.** Firms with debt access scale up quickly and adjust capital near the interior margin. Firms without access expand through a save-then-jump pattern, accumulate assets before adding plants, and can remain in a region with high marginal products of capital for a long time.

5.2 Life-cycle dynamics

We begin with simulated life-cycle paths. Figure 7 compares firms with and without debt access for a high-productivity type. The contrast is sharp. With debt access, firms scale up quickly: they add plants early, adjust capital toward target values, and smooth resources intertemporally through their bond position. Without debt access, expansion is slower and lumpy. Firms accumulate liquid assets for a period, then add a plant discretely once resources are sufficient. This save-then-jump behavior is the dynamic counterpart of the affordability region described in Section 4.

A useful feature of the figure is that the gap is not only quantitative but also qualitative. The no-access firm does not merely invest *less*; it follows a different adjustment rule. It remains constrained for longer, experiences discrete increases in scale, and can spend long stretches with a marginal product of capital above the user cost because the next desired expansion is not yet financeable.

Figure 8 shows heterogeneity across productivity types for firms without debt access. Higher-productivity firms desire larger scale, so remaining without debt access is especially costly for them. They spend longer below target scale and accumulate more resources before each discrete expansion. As a result, the no-response region is not a transient friction that vanishes mechanically after one adjustment. For some firms, especially more productive ones,

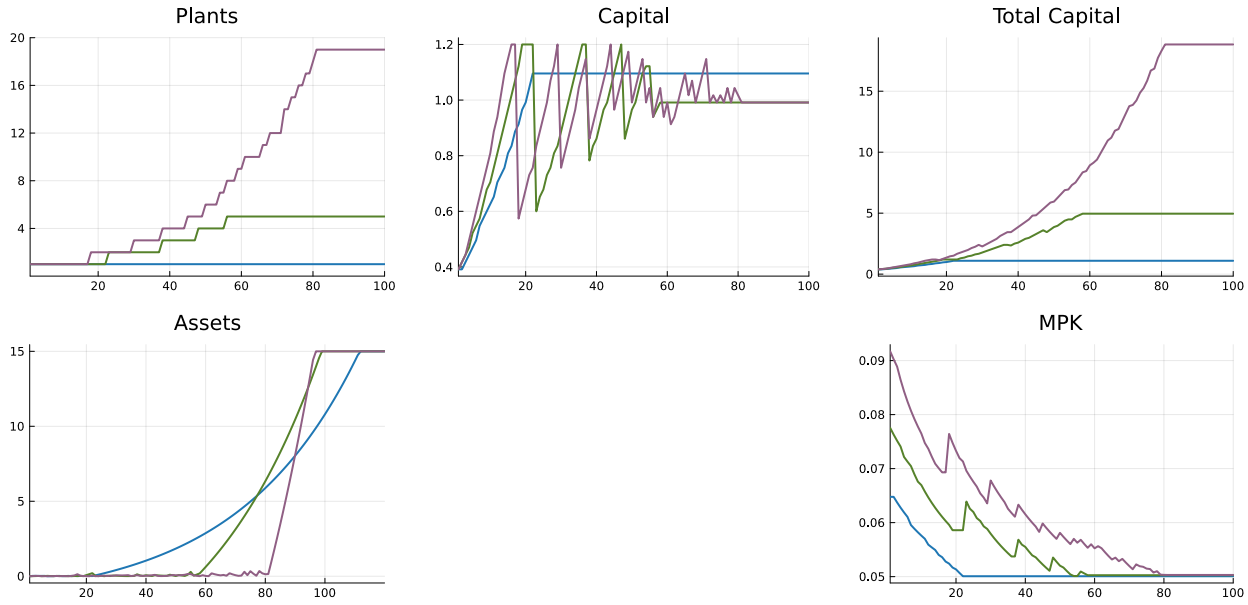


Figure 8: **Life cycle of no-access firms: heterogeneity by productivity.** More productive firms desire larger scale and therefore can remain longer in the constrained save-then-jump region when they lack debt access.

the interaction between limited access and repeated productive indivisibilities generates a persistent life-cycle phase in which monetary policy has weak impact effects on investment.

5.3 A temporary monetary policy tightening

We now study a temporary monetary policy tightening. The economy starts in the stationary equilibrium at $t = 0$. At $t = 1$, the real interest rate rises from 2.5% to 3.0% and then gradually returns to its baseline value over the next 20 periods. We solve for the perfect-foresight transition in general equilibrium, allowing the wage to adjust endogenously.

Figure 9 plots the path of the interest rate and the equilibrium wage. The increase in the interest rate lowers desired capital and labor demand, which pushes the wage down along the transition.

Figure 10 reports aggregate responses. Output falls by about 4% in period 2 and remains depressed for some time. Capital expenditures fall sharply. At the same time, plant building rises slightly and asset accumulation increases on impact. This combination is informative. Firms that are already on the interior margin reduce capital demand immediately when the user cost rises. But firms in the no-access save-then-jump region respond differently: the higher interest rate raises the return to waiting and strengthens the incentive to accumulate resources before the next discrete expansion. The aggregate response is therefore a mixture of immediate intensive-margin adjustment and delayed threshold-crossing behavior.

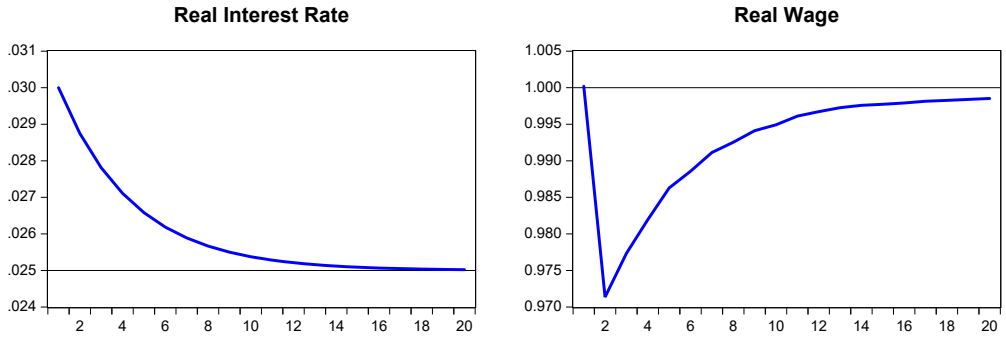


Figure 9: **Monetary policy exercise: interest rate and wage.** The real interest rate rises temporarily from 2.5% to 3.0%, and the equilibrium wage adjusts along the transition.

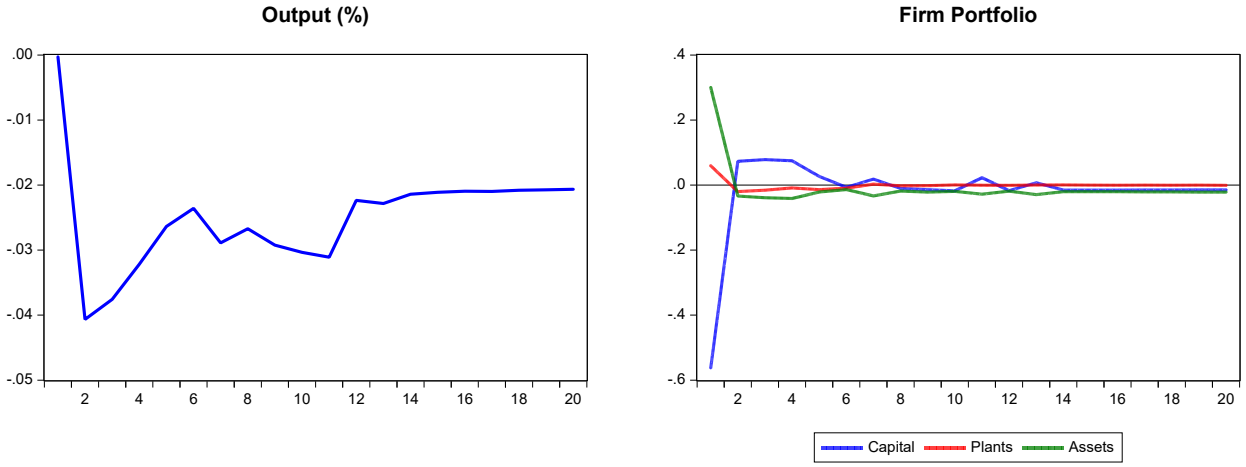


Figure 10: **Aggregate responses to a temporary monetary tightening.** Output contracts, capital expenditures fall sharply, plant creation rises slightly, and asset accumulation increases on impact.

To understand the source of these aggregate dynamics, Table 3 reports heterogeneity in impact responses across groups of firms. We classify firms by access status and whether they are already at target scale in the initial steady state. For each group, the table reports its stationary share, the fraction of firms leaving total capital unchanged on impact, the fraction leaving the number of plants unchanged, and the average impact change in total capital K, plants n, and bond holdings b'.

Three patterns stand out.

First, firms with debt access move immediately. They cut capital sharply on impact and increase asset holdings. This is the standard user-cost response in the model. These firms are already on the relevant financing margin, so the higher interest rate is transmitted directly into lower desired capital.

Table 3: Impact responses by group in the monetary policy exercise

Group	Share	Same K	Same n	ΔK	Δn	$\Delta b'$
Access, at target scale	77	0	90	-0.69	0.10	0.34
No access, at target scale	5	47	90	-0.13	0.23	0.07
No access, below target scale	17	97	97	0.00	-0.03	0.01
All firms	100	19	90	-0.56	0.06	0.30
Access	78	0	89	-0.71	0.07	0.38
No access	22	86	96	-0.03	0.03	0.02

Notes: The table reports impact responses in the temporary monetary tightening experiment. "Same K" and "Same n" denote the percentage of firms with no impact change in total capital or number of plants. The final three columns report average impact changes in total capital, plants, and next-period bond holdings. Small groups of newly accessing and newly born firms are omitted for brevity.

Second, firms without debt access and below target scale barely move on impact. For that group, 97% keep total capital unchanged and 97% keep the number of plants unchanged. Their average change in capital is essentially zero. This is the quantitative counterpart of the core empirical fact in Section 3: firms off the borrowing margin need not respond to a higher rate by changing current investment. Instead, they remain in the save-then-jump region.

Third, aggregate transmission depends on composition. The overall response is strong because most firms in the stationary distribution have debt access. But among no-access firms, the response is close to zero on impact. In that sense, the model delivers not just heterogeneous responses, but a very specific kind of heterogeneity: the effect of monetary policy depends on whether firms are in a regime where they can adjust at the margin or in one where they are waiting to cross the next lumpy financing threshold. The slightly positive average change in plants for access firms reflects composition: some access firms reduce per-plant capital immediately, while other firms that are already close to a discrete scale threshold still find it optimal to add a plant despite the tightening.

The quantitative exercise therefore reinforces the message of the empirical analysis. Monetary policy can pass through rapidly to borrowing conditions and generate sizable aggregate effects, while still leaving a large subset of firms with little or no immediate investment response. That subset is precisely the one the paper emphasizes: firms without debt access that are still saving toward the next productive expansion.

6 Conclusion

This paper studies how debt access shapes firm-level monetary transmission. Using monthly administrative data for the universe of Chilean firms, matched to a comprehensive credit registry and monetary policy surprises, we document four facts. Monetary policy surprises

pass through rapidly to the rates firms pay on new loans, with about 70% pass-through within 20 days. Investment falls with delay but with persistent effects after a contractionary surprise: a 100-basis-point surprise lowers cumulative investment by about 0.18 percentage points of assets after 24 months, roughly 3.2% of machinery-and-equipment capital. These real responses are concentrated among firms with positive debt. And investment rises by about 0.35 percentage points of assets when firms borrow for the first time and sustain debt.

We interpret these facts with a model in which firms expand through discrete plants and debt access is limited. Firms with debt access choose capital at the margin and respond to changes in the user cost in the standard way. Firms without debt access must self-finance productive expansions and often remain in a save-then-jump region in which affordability rather than the user cost determines current investment. In that region, firms can have high marginal products of capital and yet show little impact response to monetary policy.

The broader implication is that monetary transmission depends not only on interest-rate pass-through, leverage, or collateral, but also on the distribution of firms across financing regimes. In our quantitative exercise, access firms move immediately on impact while most no-access firms below target scale do not. An economy in which more firms are already on the borrowing margin will therefore react differently from one in which many firms are still saving toward the next discrete expansion. Quantifying how that composition changes over time and across episodes remains a natural next step.

A Additional empirical evidence

A.1 Leverage in the full sample

Our main result in the paper is that debt status is a first-order source of heterogeneity in monetary transmission. This does not mean leverage is irrelevant. To document that point, we estimate

$$Y^I_{j,t}(h) = \alpha_j + \alpha_{s,q(t)} + \beta_h \varepsilon_t + \beta_{h,L} \left(\varepsilon_t \times L_{jt}\right) + \Gamma X_{jt} + u_{j,t}(h), \tag{39}$$

where L_{jt} denotes leverage and the remaining notation follows the baseline local-projection specification in Section 3. Figure 11 shows that more leveraged firms are less responsive to monetary policy shocks in the full sample. This pattern is consistent with balance-sheet heterogeneity within the set of financially active firms.

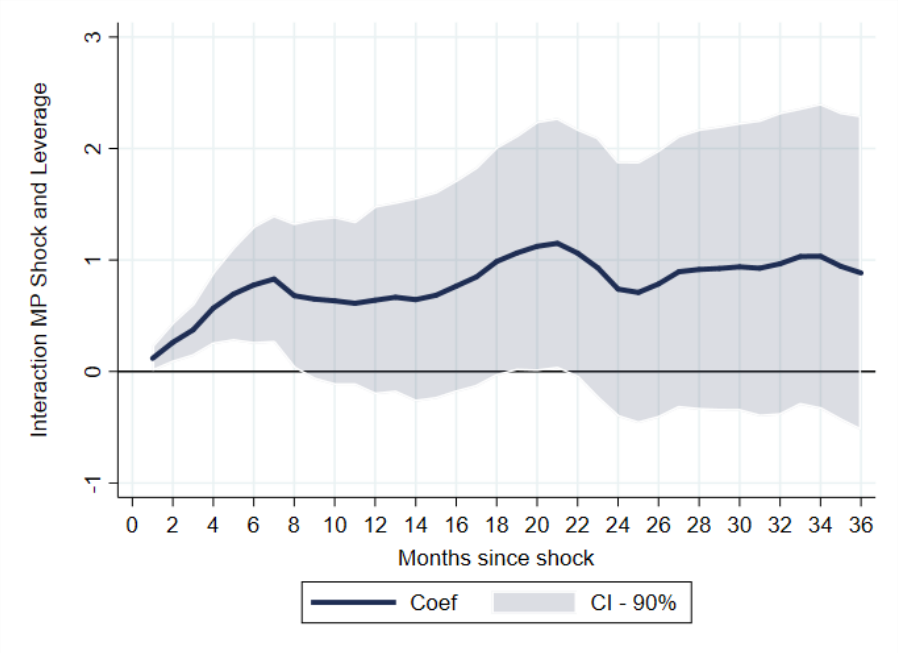


Figure 11: **Leverage interaction in the full sample.** Local projection estimates interacting the monetary policy surprise with leverage.

A.2 Public firms with debt: relation to Ottonello and Winberry (2020)

To compare our findings with Ottonello and Winberry (2020), we restrict the sample to public firms with debt and estimate the leverage interaction in that subsample. Figure 12 shows that, in this restricted sample, firms with higher leverage are consistently less responsive to the shock, in line with the Ottonello-Winberry evidence.

The comparison is useful for positioning the paper. In the relevant public-firm borrower subsample, we recover the leverage gradient emphasized by Ottonello and Winberry (2020). But once we move to the universe of firms, the more fundamental source of heterogeneity is whether firms are on the borrowing margin at all. Our debt-status result therefore complements rather than contradicts the public-firm evidence.

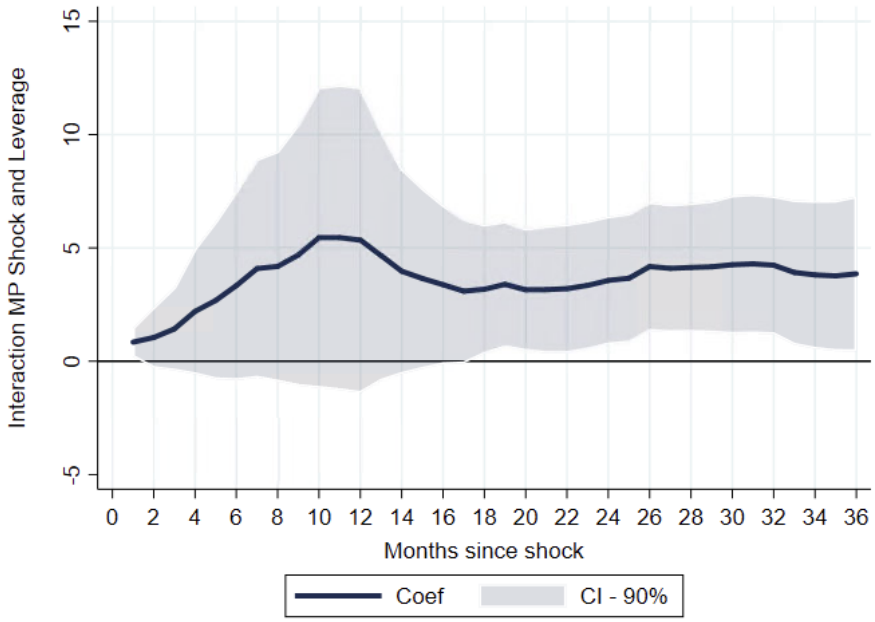


Figure 12: **Public firms with debt: Ottonello-Winberry style leverage gradient.** Restricting the sample to public firms with debt, more leveraged firms are less responsive to the monetary policy shock.

B Proofs for Section 4

B.1 Proof of Proposition 4.1

Proof. The static labor problem is

$$\max_{\ell \geq 0} \quad z n^{\xi} k^{\alpha_k} \ell^{\alpha_\ell} - w n \ell.$$

Because $\alpha_\ell \in (0,1)$, the objective is strictly concave in ℓ, so the first-order condition is necessary and sufficient:

$$\alpha_\ell z n^{\xi} k^{\alpha_k} \ell^{\alpha_\ell - 1} = w n.$$

Solving for ℓ gives

$$\ell = \left(\frac{\alpha_\ell z}{w} n^{\xi - 1} k^{\alpha_k} \right)^{\frac{1}{1-\alpha_\ell}},$$

which is (18).

Let

$$Y = z n^{\xi} k^{\alpha_k} \ell^{\alpha_\ell}.$$

Multiplying the first-order condition by ℓ gives

$$\alpha_\ell Y = wn\ell.$$

Therefore profits net of wages are

$$\pi = Y - wn\ell = (1-\alpha_\ell)Y.$$

Substituting the optimal labor choice into Y yields

$$Y = zn^{\xi}k^{\alpha_k}\left(\frac{\alpha_\ell z}{w}n^{\xi-1}k^{\alpha_k}\right)^{\frac{\alpha_\ell}{1-\alpha_\ell}}.$$

Collecting terms gives

$$Y = z^{\frac{1}{1-\alpha_\ell}}\left(\frac{\alpha_\ell}{w}\right)^{\frac{\alpha_\ell}{1-\alpha_\ell}} n^{\xi+\frac{\alpha_\ell(\xi-1)}{1-\alpha_\ell}} k^{\alpha_k+\frac{\alpha_k\alpha_\ell}{1-\alpha_\ell}}.$$

The exponents simplify to

$$\xi + \frac{\alpha_\ell(\xi-1)}{1-\alpha_\ell} = \frac{\xi-\alpha_\ell}{1-\alpha_\ell}, \qquad \alpha_k + \frac{\alpha_k\alpha_\ell}{1-\alpha_\ell} = \frac{\alpha_k}{1-\alpha_\ell}.$$

Hence

$$\pi = (1-\alpha_\ell)\, z^{\frac{1}{1-\alpha_\ell}}\left(\frac{\alpha_\ell}{w}\right)^{\frac{\alpha_\ell}{1-\alpha_\ell}} n^{\frac{\xi-\alpha_\ell}{1-\alpha_\ell}} k^{\frac{\alpha_k}{1-\alpha_\ell}},$$

which is (19). □

B.2 Proof of Lemma 4.2

Proof. From (19),

$$\pi(z,n,k;w) = C(z,w)\, n^{\frac{\xi-\alpha_\ell}{1-\alpha_\ell}} k^{\frac{\alpha_k}{1-\alpha_\ell}},$$

where

$$C(z,w) \equiv (1-\alpha_\ell)\, z^{\frac{1}{1-\alpha_\ell}}\left(\frac{\alpha_\ell}{w}\right)^{\frac{\alpha_\ell}{1-\alpha_\ell}} > 0.$$

Differentiating with respect to k,

$$\pi_k(z,n,k;w) = C(z,w)\,\frac{\alpha_k}{1-\alpha_\ell}\, n^{\frac{\xi-\alpha_\ell}{1-\alpha_\ell}} k^{\frac{\alpha_k}{1-\alpha_\ell}-1}.$$

Therefore

$$MPK(z,n,k;w) = \frac{1}{n}\pi_k(z,n,k;w) = \widetilde{C}(z,w)\, n^{\frac{\xi-1}{1-\alpha_\ell}} k^{\frac{\alpha_k+\alpha_\ell-1}{1-\alpha_\ell}},$$

where

$$\widetilde{C}(z,w) = C(z,w)\frac{\alpha_k}{1-\alpha_\ell} > 0.$$

Since $\alpha_k + \alpha_\ell < 1$, the exponent on k is negative:

$$\frac{\alpha_k + \alpha_\ell - 1}{1-\alpha_\ell} < 0.$$

Hence MPK is strictly decreasing in k and continuous on $(0, \bar{k}]$.

The exponent on n is

$$\frac{\xi - 1}{1-\alpha_\ell} \leq 0,$$

with strict inequality if $\xi < 1$. Hence MPK is weakly decreasing in n, and strictly decreasing in n if $\xi < 1$.

Continuity and strict monotonicity in k imply that whenever the equation

$$MPK(z,n,k;w) = R + \delta$$

has a solution, that solution is unique. □

B.3 A preliminary monotonicity lemma

Lemma B.1. *For fixed (z,k,n), both $V^A(z,b,k,n)$ and $V^N(z,b,k,n)$ are weakly increasing in b.*

Proof. Consider first V^A. If $b_2 > b_1$, then current cash on hand differs by

$$X(z,b_2,k,n;w,R) - X(z,b_1,k,n;w,R) = (1+R)(b_2 - b_1) > 0.$$

Any feasible plan under b_1 remains feasible under b_2, with the extra resources paid as additional dividends. Hence V^A is weakly increasing in b.

The same feasibility argument applies to V^N. In addition, higher current b raises current predetermined book value $W = b + nk + Fn$, which weakly increases the probability of waking up with access next period. Hence $V^N(z,b,k,n)$ is weakly increasing in b. □

B.4 Proof of Proposition 4.3

Proof. We prove each part in turn.

Part (i). Fix an access state (z, b, k, n). For any candidate real choice (n', k') with $n' \geq n$ and $0 \leq k' \leq \bar{k}$, the firm can always make that choice feasible by choosing sufficiently negative b' and setting $d = 0$ in the budget constraint

$$d + b' = X - n'k' - F(n' - n).$$

Hence current liquidity never restricts the set of feasible real choices. Differences in current cash on hand are absorbed by the financing decision b' and by dividends d, not by the real choice itself. Therefore the ranking of feasible real choices (n', k') is independent of current liquid wealth b and current installed capital k, which enter only through current resources.

Part (ii). Because access firms can borrow freely, any real choice that is optimal in the access regime can be implemented immediately. Delaying a scale increase serves no financing purpose. Thus, from any access state, the firm chooses the access-regime optimal real choice immediately rather than waiting to self-finance it.

Part (iii). Fix an optimal access choice with $k' \in (0, \bar{k})$. Hold n' fixed. For any sufficiently small $\eta > 0$, the perturbations

$$(k', b') \mapsto (k' + \eta,\ b' - n'\eta) \qquad \text{and} \qquad (k', b') \mapsto (k' - \eta,\ b' + n'\eta)$$

are feasible. These perturbations leave the current budget constraint unchanged because the increase in installed total capital, $n'\eta$, is exactly offset by the change in liquid wealth.

Under immediate exit after choices, liquidation value is unchanged by these perturbations:

$$(b' - n'\eta) + n'(k' + \eta) = b' + n'k',$$

and similarly for the reverse perturbation. Hence exit payoffs cancel exactly. Conditional on survival, next period's cash on hand changes by

$$n'\eta \left[MPK(z, n', k'; w) + (1 - \delta) - (1 + R) \right].$$

The term MPK captures the marginal increase in next-period operating profits, $(1 - \delta)$ is the marginal undepreciated capital carried into next period, and $(1 + R)$ is the gross return foregone on liquid wealth. If

$$MPK(z, n', k'; w) > R + \delta,$$

then the first perturbation strictly raises value. If

$$MPK(z, n', k'; w) < R + \delta,$$

then the second perturbation strictly raises value. Both contradict optimality. Therefore

$$MPK(z, n', k'; w) = R + \delta.$$

□

B.5 Proof of Proposition 4.4

Proof. We prove each part in turn.

Part (i). Suppose $b' > 0$ and $k' \in (0, \bar{k})$. Since $b' > 0$, there exists $\eta > 0$ such that decreasing b' by $n'\eta$ remains feasible. Since $k' \in (0, \bar{k})$, both increasing and decreasing k' by η are feasible while holding n' fixed.

Because current W is predetermined, the continuation weights $p(W)$ and $1 - p(W)$ are fixed at the time of choice. Consider the perturbations

$$(k', b') \mapsto (k' + \eta,\ b' - n'\eta) \qquad \text{and} \qquad (k', b') \mapsto (k' - \eta,\ b' + n'\eta)\,.$$

They leave the current budget constraint unchanged. Immediate-exit liquidation is also unchanged because

$$(b' - n'\eta) + n'(k' + \eta) = b' + n'k'.$$

Hence exit payoffs cancel exactly. Conditional on survival, next period's resources change by

$$n'\eta\left[MPK(z, n', k'; w) + (1 - \delta) - (1 + R)\right].$$

If $MPK > R + \delta$, shifting resources from liquid wealth into installed capital is profitable. If $MPK < R + \delta$, shifting resources in the other direction is profitable. Optimality therefore implies

$$MPK(z, n', k'; w) = R + \delta.$$

Part (ii). Suppose $b' = 0$. We first show that there exists an optimal representative with $d = 0$.

If $d > 0$, then by the budget constraint,

$$d = X - n'k' - F(n' - n).$$

Because $d > 0$ and $b' = 0$, choosing instead

$$\widetilde{b}' = \varepsilon, \qquad \widetilde{d} = d - \varepsilon$$

for some sufficiently small $\varepsilon \in (0, d)$ leaves (n', k') unchanged and remains feasible. By Lemma B.1, the continuation value is weakly increasing in next-period liquid wealth, so this deviation is weakly better. Hence any optimum with $b' = 0$ has an equivalent or better representative with $d = 0$.

With $b' = 0$ and $d = 0$, the budget constraint becomes

$$n'k' + F(n' - n) = X,$$

which is (31).

Part (iii). Suppose now that $b' = 0$, $k' \in (0, \bar{k})$, and the interior target for the chosen n' is not affordable. Let $\hat{k}(z, n', R; w)$ denote the unique interior target solving

$$MPK(z, n', \hat{k}; w) = R + \delta.$$

By Lemma 4.2, this target is unique whenever it exists.

Because the interior target is not affordable, the allocation

$$(n', \hat{k}, b' = 0, d = 0)$$

violates the budget constraint. Since the chosen pair (n', k') satisfies (31), affordability implies

$$n'k' < n'\hat{k}.$$

Hence

$$k' < \hat{k}(z, n', R; w).$$

Because MPK is strictly decreasing in k, it follows that

$$MPK(z, n', k'; w) > MPK(z, n', \hat{k}; w) = R + \delta.$$

This proves (32).

Part (iv). This is immediate from the definition of the plant-level capacity constraint. □

B.6 Proof of Corollary 4.5

Proof. Part (i) follows directly from the perturbation logic in Propositions 4.3 and 4.4: whenever the firm is on the interior margin with slack savings, optimality requires

$$MPK(z, n', k'; w) = R + \delta.$$

After an unanticipated increase to $R' > R$, the same argument implies that the new interior choice satisfies

$$MPK(z, n', k'; w) = R' + \delta.$$

For part (ii), if the no-access firm is in the cash-poor regime, then by Proposition 4.4 there exists an optimal representative characterized by

$$b' = 0, \qquad d = 0, \qquad n'k' + F(n' - n) = X.$$

Current investment is therefore governed by affordability. A change in the interest rate can alter the value of waiting and therefore future accumulation dynamics, but if the discrete plant choice that maximizes the constrained problem is locally unchanged, current investment need not respond on impact.

Part (iii) is immediate: if $k' = \bar{k}$, per-plant capital cannot increase further on impact. Any immediate response must occur through a change in the discrete plant choice n'. □

C Computational appendix

This appendix summarizes the numerical procedure used to solve the model. The implementation exploits the structure of the propositions in Section 4; it does not perform a brute-force search over the full control space at every state.

C.1 Stationary policies at fixed (w, R)

We discretize the state space for (b, k, n, z) and solve the access and no-access problems at fixed prices (w, R).

Access firms. Proposition 4.3 implies that, conditional on access, current liquidity affects the value level but not the ranking over real choices. The algorithm exploits this directly:

1. Enumerate feasible plant choices

$$n' \in \{n, n+1, \ldots, \bar{n}\}.$$

2. For each candidate n', recover the candidate per-plant capital from the theorem:

$$k' = k^*(z, n', R; w)$$

   if the Euler target is interior, and $k' = \bar{k}$ otherwise.

3. Recover the financing choice b' and dividends d from the budget constraint and the savings cap.

4. Evaluate the Bellman objective and select the maximizing n'.

Hence the access problem reduces to a finite search over n', rather than a brute-force search over (n', k', b').

No-access firms. For no-access firms, the algorithm exploits Proposition 4.4. For each feasible plant choice n', there are only a finite number of relevant regimes to check:

1. **Interior-saving regime:** if the Euler target $k^*(z, n', R; w)$ is affordable, set

$$k' = k^*(z, n', R; w), \qquad b' = X - n'k' - F(n' - n),$$

   with $b' > 0$.

2. **Cash-poor regime:** if the interior target is not affordable and capacity does not bind, set

$$b' = 0, \qquad d = 0, \qquad k' = \frac{X - F(n' - n)}{n'}.$$

3. **Capacity regime:** if affordability would imply $k' > \bar{k}$, set

$$b' = 0, \qquad d = 0, \qquad k' = \bar{k}.$$

Because access tomorrow depends on beginning-of-period $W = b + nk + Fn$, which is predetermined, the access probability is fixed across all candidate controls at the current

state. This further simplifies the no-access problem: the algorithm evaluates a small finite set of regime-specific candidates rather than a full search over all controls.

In implementation, knife-edge cases are assigned by tolerance. If the computed b' is numerically close to zero, the choice is classified as boundary rather than interior; if, in addition, k' is numerically close to $\bar{k}$, it is assigned to the capacity regime.

C.2 Ergodic distribution

Given stationary policies, we simulate a large panel of firms forward:

1. initialize entrants at $(a, b, k, n) = (N, b_0, k_0, n_0)$ with productivity drawn from the discrete z distribution;
2. evolve firms using the stationary policy rules;
3. apply the access lottery for no-access firms using beginning-of-period $p(W)$, with $W = b + nk + Fn$ from the pre-choice state;
4. apply iid exit with probability $1 - \zeta$;
5. replace exiting firms with new entrants.

The ergodic distribution is approximated by the invariant cross-sectional distribution generated by this simulation.

C.3 Temporary monetary policy experiment

For the perfect-foresight transition:

1. take the baseline stationary objects (w_0, R_0, μ_0) as initial conditions;
2. specify the temporary path for R_t;
3. guess a path for wages $\{w_t\}_{t=0}^{T}$;
4. solve firm problems backward given $\{R_t, w_t\}$ and terminal continuation values equal to the baseline stationary values;
5. simulate the cross-section forward from μ_0;
6. update the wage path using excess labor demand;
7. iterate until the wage path and labor-market errors converge.

This is the procedure used for the transition results in Section 5.

D Cumulative investment and log capital

This appendix records the approximation used to interpret cumulative investment scaled by assets as proportional to log changes in the capital stock. Starting from the capital accumulation identity,

$$K_{t+1} = (1-\delta)K_t + \widetilde{I}_{t+1},$$

we can write

$$\log K_{t+h} - \log K_t = \sum_{m=1}^{h} \log\left(\frac{K_{t+m}}{K_{t+m-1}}\right) = \sum_{m=1}^{h} \log\left(1 + \frac{\widetilde{I}_{t+m} - \delta K_{t+m-1}}{K_{t+m-1}}\right).$$

Using $\log(1+x) \approx x$ for small x gives

$$\log K_{t+h} - \log K_t \approx \sum_{m=1}^{h} \frac{\widetilde{I}_{t+m}}{K_{t+m-1}} - \delta h.$$

In the empirical work, we proxy the scale term with total assets and focus on cumulative machinery-and-equipment investment:

$$\sum_{m=0}^{h} \left(\frac{i_{j,t+m}}{A_{j,t+m}}\right).$$

This approximation is the basis for translating our local-projection coefficients into percentage changes in the capital stock. Since positive investment is infrequent and the investment-to-assets ratio is modest on average in our sample, the approximation is most useful as a medium-run magnitude conversion rather than as an exact accounting identity at the firm-month level.

References

Anderson, G. and Cesa-Bianchi, A. (2024). Crossing the credit channel: Credit spreads and firm heterogeneity. *American Economic Journal: Macroeconomics*, 16(3):417–446.

Aruoba, S. B., Fernández, A., Guzmán, D., Pastén, E., and Saffie, F. (2021). Monetary policy surprises in chile: Measurement and real effects. Working Paper 921, Central Bank of Chile.

Bahaj, S., Foulis, A., Pinter, G., and Surico, P. (2022). Employment and the residential collateral channel of monetary policy. *Journal of Monetary Economics*, 131:26–44.

Buera, F. J. and Shin, Y. (2013). Financial frictions and the persistence of history: A quantitative exploration. *Journal of Political Economy*, 121(2):221–272.

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

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

Durante, E., Ferrando, A., and Vermeulen, P. (2022). Monetary policy, investment and firm heterogeneity. *European Economic Review*, 148:104251.

Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.

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

Görtz, C., Sakellaris, P., and Tsoukalas, J. D. (2023). Firms' financing dynamics around lumpy capacity adjustments. *European Economic Review*, 156:104481.

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

Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.

Jeenas, P. (2024). Firm balance sheet liquidity, monetary policy shocks, and investment dynamics. Working Paper 1409, Barcelona School of Economics. Revised May 2024.

Midrigan, V. and Xu, D. Y. (2014). Finance and misallocation: Evidence from plant-level data. *American Economic Review*, 104(2):422–458.

Moll, B. (2014). Productivity losses from financial frictions: Can self-financing undo capital misallocation? *American Economic Review*, 104(10):3186–3221.

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