Monetary Policy, Firm Heterogeneity, and the Distribution of Investment Rates*

Matthias Gnewuch[1], Donghai Zhang[2,*]

Abstract

We document that an interest rate cut reshapes the cross-sectional distribution of investment rates—fewer zero and small investment rates and more large ones—and particularly so among young firms. The extensive margin investment decision—whether to invest or not—is essential in explaining these findings. We develop a heterogeneous-firm model with fixed adjustment costs and firm life-cycle dynamics to rationalize the evidence and study the implications for the investment channel. The extensive margin investment decision makes monetary policy less effective whenever few firms are inclined to invest: in downturns, but also in economies with low business dynamism and few young firms.

Keywords: Investment Rate Distribution, Adjustment Costs, Lumpy Investment, Heterogeneous Sensitivity, Extensive Margin, Monetary Policy

JEL classification: E52, E22, D21, D22

*An earlier draft was titled "On the Heterogeneous Interest Rate Sensitivity of Firms' Investments" (September 2021). We thank Isaac Baley, Christian Bayer, Yongsung Chang, Russell Cooper, Maren Froemel, Joachim Jungherr, Enisse Kharroubi, Keith Kuester, Matthias Meier, Emi Nakamura, Haozhou Tang, Petr Sedláček, Vincent Sterk as well as seminar/workshop participants at Peking University, VfS 2021, Bonn University, EEA 2022, Innsbruck Workshop in Empirical Macroeconomics 2023, T2M 2023, SED 2023, ESEM 2023, Banque de France, ENRI 2023, RES 2024, and the Ghent Workshop in Empirical Macroeconomics 2024 for insightful discussions. We thank Lixing Wang for excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily represent those of the European Stability Mechanism. Donghai Zhang thanks the National University of Singapore for financial support under the Startup grant FY2023.

*Corresponding author.

Email addresses: `matthias.gnewuch@gmail.com` (Matthias Gnewuch), `donghai.d.zhang@gmail.com` (Donghai Zhang)

[1]European Stability Mechanism, 6a Circuit de La Foire Internationale, 1347 Luxembourg, Luxembourg

[2]National University of Singapore, 1 Arts Link, Singapore 117579

1. Introduction

Understanding the investment channel of monetary policy is important for policymakers because investment is a sizable and the most volatile component of aggregate GDP. To this end, the literature has extensively studied the effect of monetary policy on the *average* investment rate (Christiano et al., 2005). However, an estimated effect on the average investment rate can reflect a *shift* of the entire distribution—if all firms increase their investment rates equally—or a change in the *shape* of the distribution. How does monetary policy affect the distribution of investment rates? Moreover, a growing literature documents heterogeneous effects of monetary policy on the investment rates of various groups of firms (e.g., Gertler and Gilchrist 1994, Ottonello and Winberry 2020, Jeenas 2023, Cloyne et al. 2023). Are these heterogeneous effects on *average* investment rates due to unequal changes in the shape of the distribution? Answers to these questions are important to understand the transmission of monetary policy. In particular, they shed light on the frictions and features of firm-level investment behavior that matter for the (heterogeneous) effects of monetary policy.

In this paper, we present three pieces of evidence that are important to understand the investment channel of monetary policy. First, monetary policy changes the *shape* of the distribution of investment rates. Specifically, an expansionary monetary policy shock leads to fewer zero and small investment rates and more large ones. Second, the change in the shape of the investment rate distribution is more pronounced among young firms than among old firms. These findings highlight the relevance of the extensive margin investment decision—whether to invest or not—for the transmission of monetary policy. Third, the extensive margin accounts for around 60% of the heterogeneous average effect of monetary policy across age groups. We rationalize these findings in a heterogeneous-firm model with fixed capital adjustment costs and life-cycle dynamics. In the presence of fixed costs, expansionary monetary policy induces some firms to switch from not investing to making a sizeable investment, which is crucial to match the effect on the distribution of investment rates. With life-cycle dynamics, fixed costs also make young firms more responsive to monetary policy—even in the absence of financial frictions—because young firms can more easily be induced to make an investment. Finally, the extensive margin investment channel makes monetary policy less effective in stimulating aggregate investment whenever only few firms are inclined to invest. This is the case in downturns, but also in economies with low business dynamism and few young and growing firms.

In more detail, we study the investment channel using quarterly firm-level investment data from Compustat in combination with identified monetary policy shocks as in Ottonello and Winberry (2020) and Cloyne et al. (2023). In contrast to the existing literature, we estimate the effects of monetary policy on the investment rate distribution—using quantiles and the share of firms in each bin of the distribution—rather than solely focusing on the average investment rate. We show that after an expansionary monetary policy shock, there are fewer firms making a small or no investment and more firms making a large investment, altering the shape of the distribution of investment rates—**Fact 1**. This novel evidence highlights the presence of a quantitatively relevant investment channel of monetary policy along the extensive margin.

Conducting the same empirical analysis for young and old firms separately, we uncover that the effect of monetary policy on the shape of the distribution of investment rates is more pronounced among young firms than among old firms—**Fact 2**. This finding suggests that the extensive margin investment channel is particularly important for young firms. We corroborate this view by showing that monetary policy increases the share of investment spikes ($i > 10\%$) and decreases the share of inaction ($|i| < 0.5\%$) more strongly for young firms than for old firms. These findings are present even among firms deemed unlikely to be financially constrained—characterized by low leverage, high liquidity, or having paid dividends—supporting the interpretation that these heterogeneous effects do not reflect financial frictions. Finally, we estimate the extensive margin—using the spike rate as a proxy—to account for around 60% of the heterogeneous effect of monetary policy on average investment rates of young and old firms—**Fact 3**.

The second part of the paper interprets the empirical findings through the lens of a general equilibrium heterogeneous-firm model with fixed capital adjustment costs, firm life-cycle dynamics, and nominal rigidities. The model, calibrated to aggregate and firm-level data from the U.S., replicates all three empirical facts. Fixed costs give rise to lumpy investment behavior and an investment channel of monetary policy along the extensive margin. After an expansionary monetary policy shock, some firms switch from being inactive to making a sizeable investment, leading to a change in the distribution of investment rates consistent with **Fact 1**. Entry and exit give rise to firm life-cycle profiles and an age distribution. Young firms are on average farther away from their optimal level of capital than old firms, and can therefore more easily be induced to make an investment. In consequence, monetary policy has a heterogeneous effect on spike rates, average investment rates, and distributions across age groups, consistent with **Fact 2**. The heterogeneous average effect is predominantly explained by the extensive margin investment decision, consistent with **Fact 3**.

Next, we explore the aggregate implications of the heterogeneous-firm model with fixed adjustment costs and firm life-cycle dynamics. Broadly speaking, the model implies that monetary policy is particularly effective whenever there are many firms that can easily be induced to make a large investment. We provide two examples to emphasize that both long-run trends and cyclical developments are quantitatively relevant. First, we show that the decline in firm dynamism and ensuing "aging" of the firm distribution (i.e., lower share of young firms) observed since the 1980s has, according to the model, made monetary policy about 12% less effective in stimulating investment. Second, consistent with Winberry (2021) and Koby and Wolf (2020), monetary policy is less effective in a recession than in a boom, because in a recession, fewer firms are inclined to invest. The latter implication aligns well with the empirical literature showing that monetary policy interventions are less potent in recessions (Tenreyro and Thwaites, 2016).

Finally, our three empirical facts, paired with insights from the heterogeneous-firm model with lumpy investment, offer several implications that are helpful to understand the investment channel of monetary policy. On the one hand, our analysis highlights the importance of the extensive margin investment decision and fixed adjustment costs. Monetary policy stimulates aggregate investment not because many firms increase their investment rates a little, but rather because few firms switch from inaction to making a sizeable investment, thus increasing their investment rates by a lot. On the other hand, our analysis offers insights about financial frictions and financial acceleration. This is because young firm age is oftentimes used as a proxy variable for tight financial constraints and a heterogeneous effect of monetary policy across age groups is used as evidence for financial acceleration (Cloyne et al., 2023; Gertler and Gilchrist, 1994). Our model shows that financial acceleration is not necessary to generate a heterogeneous effect across age groups. Hence, there is an issue of observational equivalence as we show that a model with a non-financial friction (fixed adjustment costs) equally generates a heterogeneous effect across age groups. The difference between both frictions is important for policymaking because financial acceleration implies that the effectiveness of monetary policy is countercyclical—more firms are financially constrained in recessions—whereas fixed costs imply a procyclical effectiveness.

Literature Review. First and foremost, this paper relates to the empirical literature studying the investment channel of monetary policy using aggregate data (e.g., Christiano et al. 2005), and firm-level data (e.g., Gertler and Gilchrist 1994, Ottonello and Winberry 2020, Jeenas 2023, Cloyne et al. 2023, Jungherr et al. 2022, González et al. 2022). So far, this literature has focused on the effects on aggregate investment or average investment rates. Our first contribution is to document how monetary policy affects the entire distribution of investment rates as well as moments thereof, such as the spike rate and the inaction rate.[3] A strand of this literature studies the heterogeneous effects across various groups of firms. We, secondly, contribute to this strand by showing that between young and old firms, not only the average effects differ, but also the effects on the distribution as well as on spike and inaction rates. Both of these empirical contributions are important to understand which features of firm-level investment behavior—lumpy investment, life-cycle dynamics, and financial frictions—matter for the transmission of monetary policy.

Second, this paper relates to the theoretical and quantitative literature studying the relevance of the extensive margin investment decision and lumpy investment behavior for aggregate investment, particularly for its responsiveness to shocks over the business cycle. Important contributions include Caballero et al. (1995), Caballero and Engel (1999), Thomas (2002), Khan and Thomas (2003), Khan and Thomas (2008), Bachmann et al. (2013), Koby and Wolf (2020), Winberry (2021), and Baley and Blanco (2021). Monetary policy shocks

[3]Gourio and Kashyap (2007) estimate the unconditional cyclicality of the spike rate of firms' investments but do not investigate monetary policy shocks. In contemporaneous work, Lee (2023) estimates the effect of monetary policy shocks on the spike rates of small and large firms but neither studies young and old firms nor the entire distribution.

in models with fixed adjustment costs have been analyzed by Reiter et al. (2013), Reiter et al. (2020), and Fang (2023). We contribute to this literature by incorporating firm life cycles into an otherwise standard heterogeneous-firm model with lumpy investment. The combination of these two features is important for two reasons. First, introducing firm life cycles allows us to examine heterogeneous distributional effects of monetary policy across firm age groups, which is necessary to rationalize our empirical findings. Second, the combination of lumpy investment and life cycles is not only important for the cross-section but also for the *aggregate* investment channel of monetary policy. Specifically, we quantify the weakening of the effectiveness of monetary policy due to the decline in business dynamism.

The remainder of this paper is organized as follows. Section 2 presents our empirical results. Section 3 describes the New Keynesian heterogeneous-firm model. Section 4 calibrates the model and analyzes the effects of monetary policy. Section 5 concludes.

2. Empirical Evidence

We present three pieces of evidence that are important to understand the investment channel of monetary policy. After describing the data (Section 2.1) and local projection model (Section 2.2), we document the effects of monetary policy on the distribution of investment rates (Section 2.3). Then, we present the heterogeneous effects of monetary policy by firm age (Section 2.4) and estimate the contribution of the extensive margin to heterogeneous average effects (Section 2.5). In addition, we discuss how our findings relate to the work of Cloyne et al. (2023) (Section 2.6).

2.1. Firm-Level Data

We use quarterly firm-level data from Compustat. Our sample begins with 1986Q1 and ends with 2018Q4. We exclude firms located outside the United States, with incomplete or questionable information (e.g. negative reported sales) and those not suitable for our analysis (e.g. financial firms) from the sample. Details on the sample selection are relegated to Appendix D.1. Since information on firm age in Compustat is scarce, we merge age information from WorldScope and Jay Ritter's database, as explained in Appendix D.2.

Capital stocks reported in Compustat are accounting capital stocks and do not perfectly reflect economic capital stocks. To address this issue, we use a Perpetual Inventory Method (PIM) to compute real economic capital stocks, building on Bachmann and Bayer (2014). Details of this procedure are explained in Appendix D.3. Our baseline measure of the investment rate of firm j at time t is $i_{jt} = \frac{CAPX_{jt} - SPPE_{jt}}{INVDEF_t \times k_{jt-1}}$, thus, real capital expenditures (CAPX) net of sales of capital (SPPE) divided by the lagged real economic capital stock (k). More details on the construction of variables are given in Appendix D.4.

For parts of the subsequent analysis, we aggregate the firm-level data to quarterly investment rate distributions and moments thereof.[4] The distribution of investment rates, shown in panel (c) of Figure 1, depicts some well-known features of investment rate distributions. That is, the distribution has a positive skewness, a long right tail, substantial mass at 0, and very few negative observations.

2.2. Local Projections

To estimate the effects of monetary policy shocks, we estimate the following simple local projection (LP) models:

$$y_{t+h} - y_{t-1} = \alpha^h + \beta^h \epsilon_t^{MP} + \sum_{j=2}^{4} \gamma^j \mathbb{1}\{q_{t+h} = j\} + e_{t+h} \tag{1}$$

where y_t indicates the outcome variable, ϵ_t^{MP} is the monetary policy shock, q_t is the calendar quarter, and $\mathbb{1}\{q_{t+h} = j\}$ are quarter dummies that are included to address seasonality. We use the monetary policy shocks implied by the Proxy SVAR in Gertler and Karadi (2015). These are extracted after updating the time series data used in the VAR as well as the high-frequency instruments. Details are relegated to Appendix D.5. Unless stated otherwise, the shocks are scaled to reduce the 1-year Treasury yield on impact by 25 basis points. Throughout, we use Newey-West standard errors to account for heteroskedasticity and autocorrelation. Before turning to our novel findings, we verify that the monetary policy shocks have plausible effects on aggregate variables. We show in Appendix D.6 that an expansionary shock leads to hump-shaped increases in both investment and GDP. The peak effects are 1.4% (investment) and 0.35% (GDP), respectively.

[4]Moments which are sensitive to outliers (e.g., mean), are winsorized by group and quarter.

2.3. Fact 1: Shape of the Distribution of Investment Rates

The literature has extensively studied the effect of monetary policy on the *average* investment rate.[5] On the one hand, this estimated effect on the average investment rate could reflect that all firms increase their investment rate by the same (average) amount. In this case, we would expect the distribution of investment rates to *shift* to the right, but not change its shape. On the other hand, the change in the average investment rate could reflect that only a few firms increase their investment rate, but by a large amount. In this case, we would expect to see a change in the *shape* of the distribution of investment rates.

To investigate whether monetary policy affects the distribution of firm-level investment rates, we estimate the effects on different quantiles of the investment rate distribution, using the time series of the respective quantiles as outcome variables in the empirical model (equation 1).[6] If the increase in the average investment rate reflects a mere shifting of the distribution, the effect on all quantiles must be identical. Panel (a) of Figure 1 shows the effect of a monetary policy shock on quantiles of the investment rate distribution. It is evident that the right tail of the investment rate distribution (75th percentile) responds more strongly than the left tail (25th percentile). This difference is statistically significant, as illustrated by the effect on the corresponding interquartile range in panel (b). These findings are robust to alternative choices of quantiles as shown in Figure A.1. The disproportionate change in the right tail compared to the left tail shows that monetary policy changes the shape of the investment rate distribution.

To investigate in more detail how monetary policy affects the distribution of investment rates, we use a binned distribution and regress the time series of the *share* of firms in each bin on the monetary policy shock. Panel (c) of Figure 1 shows the average distribution of investment rates (blue bars) next to the predicted distribution at the horizon at which the effect of the monetary policy shock peaks (red bars).[7] Panel (b) illustrates the difference between the two distributions. Confirming the evidence from panels (a) and (b), there is

[5] We show the effect of monetary policy on the average investment rate in panel (a) of Figure 4.

[6] Loria et al. (2023) have recently applied a similar two-step quantile local projection approach to estimate the effects of macroeconomic shocks on the conditional quantiles of GDP growth.

[7] Horizon 13 is when the effect on the average investment rate peaks (see panel (a) of Figure 4).

Figure 1: Effect of Monetary Policy on the Distribution of Investment Rates

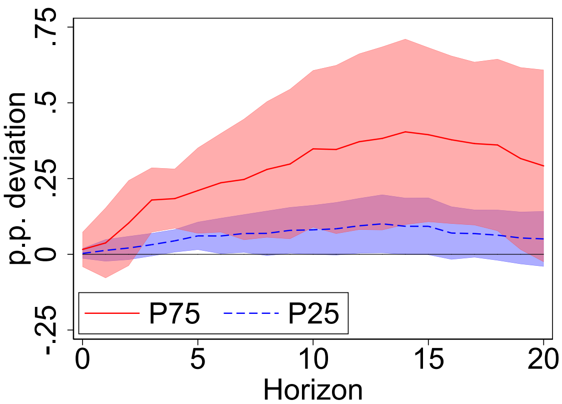

(a) Quantiles (75th & 25th)

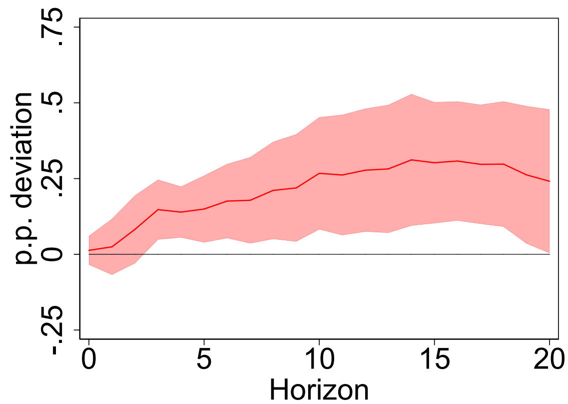

(b) Interquartile Range (P75 - P25)

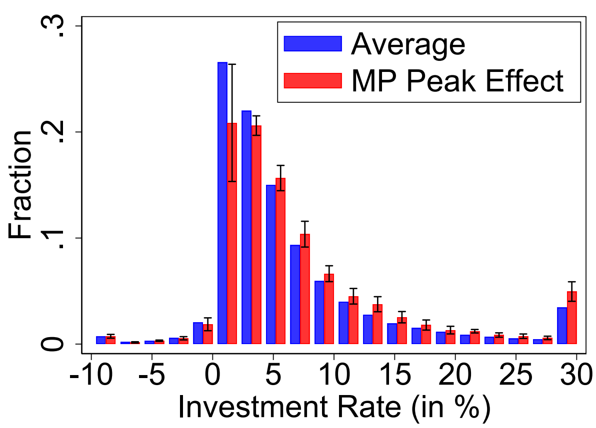

(c) Distribution (Avg. & After MP Shock)

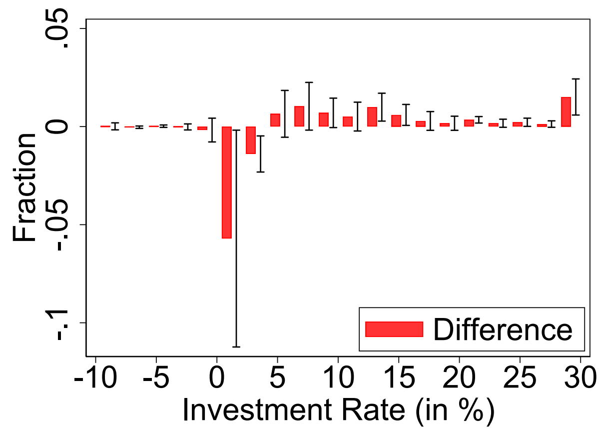

(d) Change in Distribution

Notes: Panels (a) and (b) plot the effect of a monetary policy shock on quantiles and the interquartile range of the distribution of investment rates. The lines represent the estimated $\hat{\beta}^h$ from separate regressions: $y_{t+h} - y_{t-1} = \alpha^h + \beta^h \epsilon_t^{MP} + \sum_{j=2}^{4} \gamma^j \mathbb{1}\{q_{t+h} = j\} + e_{t+h}$. The monetary policy shocks are scaled to reduce the 1-year Treasury yield by 25 basis points. The shaded areas indicate the 90% confidence intervals constructed using standard errors that are robust to heteroskedasticity and autocorrelation. Panel (c) plots the effect of a monetary policy shock on bins of the investment rate distribution. Blue bars depict the average distribution, red bars depict the predicted distribution at horizon 13 (peak effect) after a monetary policy shock. Panel (d) plots the difference between the bars in panel (c). Black lines indicate the 90% confidence intervals constructed using standard errors that are robust to heteroskedasticity and autocorrelation. To improve readability, the shocks are scaled to reduce the 1-year Treasury yield by 100 basis points in panels (c) and (d).

a marked change in the distribution of investment rates—**Fact 1**. In particular, after an expansionary monetary policy shock, there are fewer zero and small investment rates and more large investment rates. The share of firms in the bins [0, 2) and [2, 4) falls significantly, while the share of firms in all other positive bins rises, most sizably and significantly in the bin with the largest investment rates (bin $[28, \infty)$). In contrast, the share and distribution of negative investment rates are not meaningfully affected. These findings suggest that the effect of monetary policy on the *average* investment rate is driven to a sizable degree by the extensive margin, i.e., a few firms switch from making a small or no investment to making a large investment.

Effects on the Spike Rate and Inaction Rate. To further investigate the interpretation that the extensive margin investment decision is important for the effect of monetary policy on firm investment behavior, we look at two additional statistics of the investment rate distribution. These are the *spike rate*, defined as the fraction of firms whose quarterly investment rate exceeds 10%, and the *inaction rate*, defined as the fraction of firms whose quarterly investment rate is smaller than 0.5% in absolute value.[8] Corroborating our interpretation, we find that following an expansionary monetary policy shock, the inaction rate falls and the spike rate rises, as shown in Figure A.2.

2.4. Fact 2: Heterogeneous Effects across Age Groups

Our empirical strategy can also be applied to investigate *group-specific* investment rate distributions. Cloyne et al. (2023) have documented that after an expansionary monetary policy shock, young firms increase their investment rates *on average* by much more than

[8]In annual data, an investment spike is typically defined as an investment rate above 20%, so about twice the average investment rate, which, in most representative datasets, ranges between 10% and 12% (Zwick and Mahon, 2017). Since we do not use annual, but quarterly data and Compustat features higher average investment rates (shown in Figure E.2), we define an investment spike to be a quarterly investment rate exceeding 10%. This too is an investment rate roughly twice the average investment rate. In Figure A.8 in the Appendix, we show that our main results are robust to using alternative thresholds (8%, 12%) for investment spikes. Inaction is typically defined as an annual investment rate less than 1% in absolute value. For the same reasons as above, we define inaction as a quarterly investment rate smaller than 0.5% in absolute value.

old firms. We replicate this finding in Figure A.3. This difference in average effects could reflect the intensive margin—young firms changing their investment rates by more than old firms—or the extensive margin—more young firms than old firms changing their decision whether to invest at all. To understand the role of the extensive margin in explaining the existing evidence, we estimate the effect of monetary policy on age-specific investment rate distributions.

Heterogeneous Effect on Investment Rate Distributions. Again, we resort to binned distributions to investigate the effect of monetary policy on age-specific investment rate distributions. Figure 2 compares the average distribution of investment rates of young and old firms with the predicted distributions after a monetary policy shock. We find that the shape of the distribution changes more sizably and significantly for young firms—**Fact 2**. In particular, the decrease in zero and small investment rates (bin $[0, 2)$) and the increase in very large investment rates (bin $[28, \infty)$) are more pronounced and statistically significant. This suggests that the extensive margin is important to understand not only the *average* effect of monetary policy on investment rates but also the *heterogeneous* average effect across age groups.

Figures A.4 and A.5 show that the disproportionate effects of monetary policy on the right tail of the investment rate distribution (i.e., the upper quantiles), documented for all firms in Figure 1, are present among both the group of young firms and the group of old firms. However, these effects are quantitatively much more pronounced among young firms, in line with the effects on the respective distributions.

Heterogeneous Effects on Spike Rates and Inaction Rates. To lend further support to the interpretation that the extensive margin is important for the heterogeneous responsiveness of young and old firms, we look at two additional statistics of the investment rate distribution—the spike rate and the inaction rate. Figure 3 shows that the spike rate rises and the inaction rate drops more strongly for young firms, both differences being statistically significant.[9]

[9] Appendix C shows that similar but quantitatively less pronounced findings emerge when we compare small and large firms, instead of young and old firms.

Figure 2: Effect of Monetary Policy on Age-Specific Distributions of Investment Rates

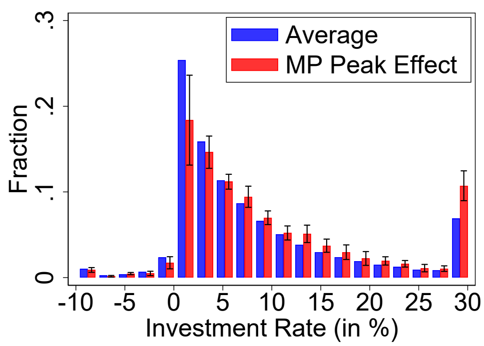

(a) Young Firms (Distributions)

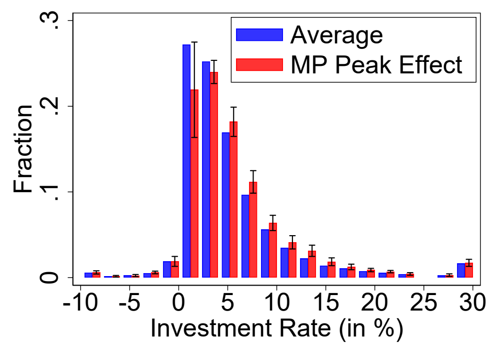

(b) Old Firms (Distributions)

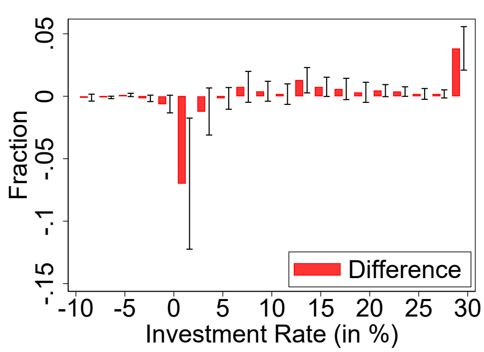

(c) Young Firms (Difference)

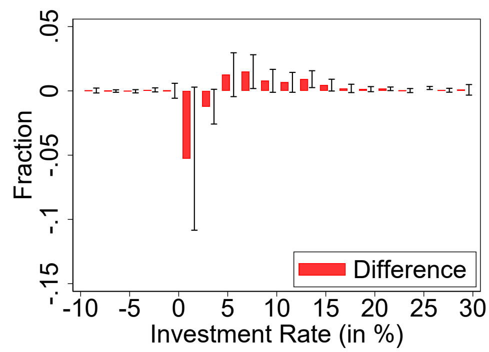

(d) Old Firms (Difference)

Notes: Panels (a) and (b) plot the effect of a monetary policy shock on bins of the investment rate distribution for young (a) and old (b) firms. Young (old) firms are less (more) than 15 years old. Blue bars depict the average distribution, red bars depict the predicted distribution at horizon 13 (peak effect) after a monetary policy shock. Panels (c) and (d) plot the difference between the bars in panels (a) and (b). Black lines indicate the 90% confidence intervals constructed using standard errors that are robust to heteroskedasticity and autocorrelation. To improve readability, the shocks are scaled to reduce the 1-year Treasury yield by 100 basis points.

Figure 3: Effect of Monetary Policy on Age-Specific Spike & Inaction Rates

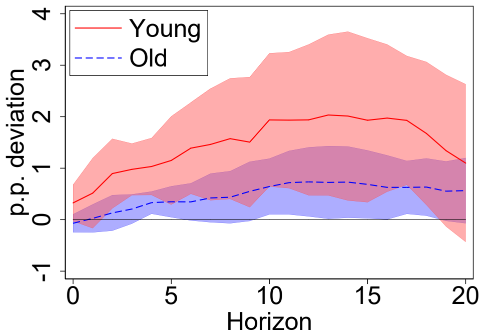

(a) Spike Rate

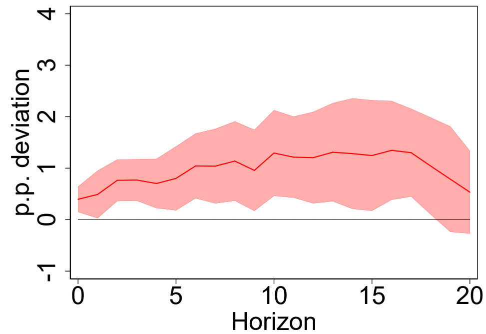

(b) Difference (Young - Old)

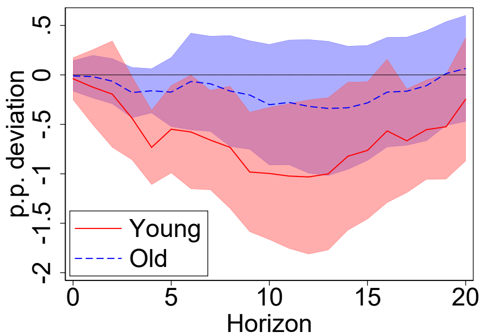

(c) Inaction Rate

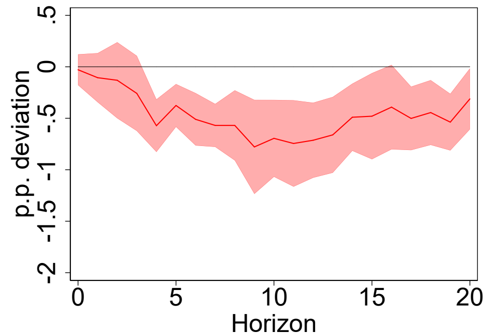

(d) Difference (Young - Old)

Notes: This figure plots the effect of a monetary policy shock on the spike rate and the inaction rate of young and old firms. Young (old) firms are less (more) than 15 years old. A spike is an investment rate exceeding 10%, inaction is an investment rate less than 0.5% in absolute value. The lines represent the estimated $\hat{\beta^h}$ from separate regressions: $y_{t+h} - y_{t-1} = \alpha^h + \beta^h \epsilon_t^{MP} + \sum_{j=2}^{4} \gamma^j \mathbb{1}\{q_{t+h} = j\} + e_{t+h}$. The monetary policy shocks are scaled to reduce the 1-year Treasury yield by 25 basis points. The shaded areas are the 90% confidence intervals constructed using standard errors that are robust to heteroskedasticity and autocorrelation.

2.5. Fact 3: The Relative Importance of the Extensive Margin

Finally, we perform a decomposition exercise to quantify the relative importance of the intensive and extensive margin, both for the effect of monetary policy on the average investment rate and the heterogeneous average effect across age groups. For this purpose, we classify investment rate observations into "spikes" ($i > 10\%$, as before) and "normal" investments ($i \leq 10\%$).[10] It follows that the average (potentially group-specific) investment rate in period t is

$$\overline{i_t} = \psi_t i_t^s + (1 - \psi_t) i_t^n \tag{2}$$

where ψ_t is the fraction of firms undertaking a "spike" in period t, and i_t^s and i_t^n are the average investment rates conditional on "spike" and "normal", respectively.[11] Then, the effect of a monetary policy shock on the average investment rate can be decomposed as follows:[12]

$$\frac{\partial \mathbb{E}(\overline{i_t})}{\partial \epsilon^{MP}} \approx \underbrace{\frac{\partial \mathbb{E}(\psi_t)}{\partial \epsilon^{MP}} \left(\mathbb{E}(i_t^s) - \mathbb{E}(i_t^n)\right)}_{\text{Extensive Margin}} + \underbrace{\mathbb{E}(\psi_t) \frac{\partial \mathbb{E}(i_t^s)}{\partial \epsilon^{MP}} + (1 - \mathbb{E}(\psi_t)) \frac{\partial \mathbb{E}(i_t^n)}{\partial \epsilon^{MP}}}_{\text{Intensive Margin}} \tag{3}$$

Intuitively, the extensive margin component reflects the change in the average investment rate that results *only* from changes in the spike rate. Vice versa, the intensive margin component reflects the change in the average investment rate that results *only* from changes in the conditional investment rates.

Panel (a) of Figure 4 plots the effect of monetary policy on the average investment rate as well as the decomposition. The extensive (intensive) margin contributes, on average across horizons, around 60% (40%) to the effect of monetary policy on the average investment rate. This finding aligns well with the evidence that *unconditional* fluctuations in aggregate investment are driven primarily by the extensive margin (Gourio and Kashyap, 2007).[13]

Decomposition of the Heterogeneous Effect of Monetary Policy. Panel (b) of Figure 4 plots the effect of an expansionary monetary policy shock on the difference between the average

[10] Figure A.8 shows that the results are robust to alternative thresholds (8%, 12%) for investment spikes.

[11] In our data, the share of (large) negative investment rates is very small and not significantly affected by monetary policy shocks. Therefore, we focus on positive spikes and disregard negative spikes.

[12] This decomposition ignores two covariance terms ($Cov(\psi_t, i_t^s)$, $Cov(\psi_t, i_t^n)$) because their contribution to the total effect is negligible. The implementation of the decomposition is explained in Appendix D.7.

[13] We find the contribution of the extensive (intensive) margin to *unconditional* fluctuations in the average investment rate to be around 60% (40%), as estimated in Appendix D.8. Hence, the contribution of the extensive margin to conditional (on monetary policy shocks) and unconditional fluctuations is similar. In Appendix E.4, we explain under which circumstances the extensive margin contributes more or less to fluctuations in the average investment rate and what implications can be drawn from the comparison of these empirical findings for the investment channel of monetary policy.

Figure 4: Decomposition of the Average Effect of Monetary Policy: Ext. & Int. Margin

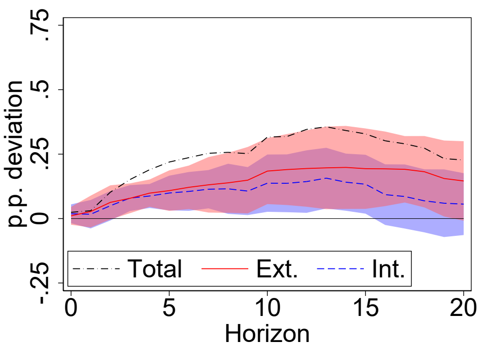

(a) Average Investment Rate (All Firms)

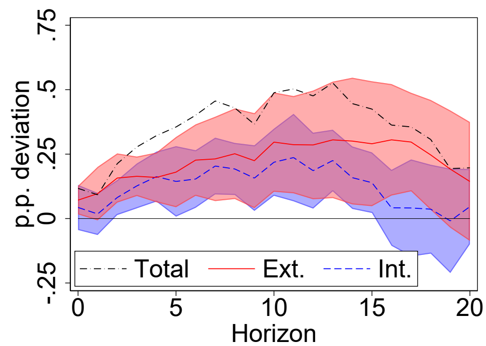

(b) Heterogeneous Effect (Young - Old)

Notes: Panel (a) shows the effect of a monetary policy shock on the average investment rate and decomposes this effect into an intensive and an extensive margin contribution, using equation (3). Panel (b) shows the heterogeneous effect of a monetary policy shock on the average investment rate of young firms as opposed to old firms and decomposes this heterogeneous effect into an intensive and an extensive margin contribution. The lines represent the estimated $\hat{\beta}^h$ from separate regressions: $y_{t+h} - y_{t-1} = \alpha^h + \beta^h \epsilon_t^{MP} + \sum_{j=2}^{4} \gamma^j \mathbb{1}\{q_{t+h} = j\} + e_{t+h}$. The monetary policy shocks are scaled to reduce the 1-year Treasury yield by 25 basis points. The shaded areas are the 90% confidence intervals constructed using standard errors that are robust to heteroskedasticity and autocorrelation.

investment rates of young and old firms. The positive effect reflects that the average investment rate of young firms is more responsive to monetary policy, confirming the finding of Cloyne et al. (2023). Panel (b) of Figure 4 also decomposes the heterogeneous effect, showing that the extensive (intensive) margin explains, on average across horizons, around 60% (40%) of the heterogeneous average effect on young and old firms—**Fact 3**.

2.6. The Role of Financial Frictions: Comparison with Cloyne et al. (2023)

Cloyne et al. (2023) interpret the finding that young firms are particularly responsive to monetary policy along the lines of the financial accelerator mechanism. Firm age is commonly used to proxy for financial constraints, reflecting the hypothesis that access to external finance depends on collateral or reputation, which young firms may lack. When monetary policy increases collateral values, this facilitates access to finance for financially constrained firms, allowing them to increase investment disproportionally. Cloyne et al. (2023) support this interpretation of the evidence by showing that the sensitivity to monetary policy is particularly strong among young firms that additionally do not pay dividends—another proxy for financial constraints—and that collateral values move procyclically after a monetary policy shock.

Our findings are consistent with the evidence provided in Cloyne et al. (2023) and are not necessarily at odds with the financial accelerator mechanism or the hypothesis that young and old firms differ in their access to finance. However, we provide *additional* evidence that suggests the presence of a *complementary* non-financial mechanism making young firms more sensitive to monetary policy. This highlights that next to financial constraints, there are additional important differences between young and old firms.[14] In particular, we show that young firms have higher average investment rates, invest more frequently (both shown in Figure 5), and are more likely to change their extensive margin investment decision in response to monetary policy (Figure 3).

[14]This argument aligns well with Crouzet and Mehrotra (2020) who argue that large firms are less cyclical than small firms because they are better diversified across industries, but not because of financial frictions, and Farre-Mensa and Ljungqvist (2016) who show that firms classified as financially constrained differ from unconstrained firms also along non-financial dimensions.

At least two pieces of evidence suggest that our mechanism—young firms being more sensitive along the extensive margin—does not reflect financial acceleration. First, we replicate the evidence presented in Figure 3 among groups of firms that are unlikely to be financially constrained. Figure A.6 shows that even among firms that have low leverage, high liquidity, or have paid dividends, the spike and inaction rates of young firms are more responsive to monetary policy shocks than those of old firms, corroborating that the heterogeneous effects of monetary policy along the extensive margin do not reflect financial acceleration. Figure A.7 shows that this result is equally present among firms that are likely financially constrained. Second, Cloyne et al. (2023) show that young firms have on average a higher Tobin's Q (their Figure 1) and that firms with a high Tobin's Q are more responsive to monetary policy (their Figure E.2), also within the subgroup of young and non-dividend-paying firms (their Figure F.5). This is in line with our interpretation that firms with high growth opportunities—such as young firms—are particularly sensitive to monetary policy because they can more easily be induced to make an investment.

Yet, the extensive margin investment channel cannot rationalize the entire difference in responsiveness to monetary policy between young and old firms. Figure 4 suggests that both the extensive and the intensive margin are important. The quantitatively relevant role of the intensive margin investment channel, which is also stronger among young firms, may very well reflect financial acceleration, in line with the interpretation of Cloyne et al. (2023).

3. Model

We build a New Keynesian model with heterogeneous firms subject to fixed and convex capital adjustment costs in the spirit of Khan and Thomas (2008) and Winberry (2021). We introduce firm entry and exit, and consequently, firm life cycles, allowing us to study not only the aggregate effects of monetary policy, but also the effects on age-specific investment rate distributions.

3.1. Investment Block

There exists a unit mass of production firms in the economy. Each firm j produces a quantity y_{jt} of the intermediate good using the production function

$$y_{jt} = z_{jt} k_{jt}^{\theta} n_{jt}^{\nu} \qquad \text{with } \theta, \nu > 0 \text{ and } \theta + \nu < 1 \tag{4}$$

where z_{jt} is total factor productivity (TFP), k_{jt} is the capital stock, and n_{jt} is the labor input. Productivity z_{jt} is subject to idiosyncratic shocks and follows an AR(1) process in logs

$$\log z_{jt} = \rho_z \log z_{jt-1} + \sigma_z \epsilon_{jt}^{z} \qquad \text{with } \epsilon_{jt}^{z} \sim \mathcal{N}(0,1) \tag{5}$$

Labor n_{jt} can be adjusted frictionlessly in every period. Capital k_{jt} is accumulated according to

$$k_{jt+1} = (1-\delta) k_{jt} + i_{jt} \tag{6}$$

where i_{jt} is investment and δ the depreciation rate.

Following Bachmann et al. (2013), we include maintenance investment. That is, a fraction χ of the depreciation δk_{jt} that occurs during the production process needs to be replaced immediately. At the end of the period, firms have $(1-\delta(1-\chi))k_{jt}$ units of capital and decide how much to invest voluntarily. To this voluntary investment, i_{jt}^{v}, there are capital adjustment costs, which need to be paid if $i_{jt}^{v} \neq 0$.[15] Total adjustment costs consist of a random fixed adjustment cost $w_t \xi_{jt}$, where ξ_{jt} is distributed uniformly between 0 and $\bar{\xi}$, and a convex adjustment cost $\frac{\phi}{2} \frac{(i_{jt}^{v})^2}{k_{jt}}$:

$$AC(k_{jt}, k_{jt+1}, \xi_{jt}) = w_t \xi_{jt} \mathbb{1}\{k_{jt+1} \neq (1-\delta(1-\chi))k_{jt}\} + \frac{\phi}{2} \frac{(k_{jt+1} - (1-\delta(1-\chi))k_{jt})^2}{k_{jt}} \tag{7}$$

where w_t is the real wage. Total investment is the sum of voluntary investment and maintenance investment. The relative price of capital (in terms of the final good) is q_t.

Entry & Exit. Firms face independent and identically distributed (i.i.d.) exit shocks ϵ_{jt}^{exit} and are forced to exit the economy at the end of the period with probability π^{exit}. Each

[15] Matching the empirical distribution of investment rates requires a rich adjustment cost specification, as discussed in Cooper and Haltiwanger (2006).

period, a fixed mass of newborn firms enters the economy. These entrants are endowed with k_0 units of capital and draw their initial (log) productivity level from the distribution $\mu^{ent} \sim \mathcal{N}(0, \frac{\sigma_z^2}{1-\rho_z^2})$, which is the ergodic distribution of equation (5).

Timing. Within any period, the timing is as follows. At stage one, idiosyncratic TFP shocks to incumbent firms realize. At stage two, a fixed mass of firms enters the economy. Entrants draw their initial productivity from μ^{ent} and are endowed with k_0 units of capital from the household. Henceforth, they are indistinguishable from incumbent firms. At stage three, firms hire labor and production takes place. Firms conduct maintenance investment. At stage four, exit shocks realize and random fixed adjustment costs are drawn. Exiting firms sell their capital stock and leave the economy. Continuing firms decide whether to adjust their capital stock or remain inactive.

Value Functions. We characterize the firm optimization problem recursively. The individual state variables are total factor productivity z and capital k. Subscripts for individual variables are henceforth dropped for readability and primes denote next period's values. The beginning-of-period real firm value is

$$V_t(z,k) = \max_n p_t z k^\theta n^\nu - w_t n + \pi^{exit} CV_t^{exit}(z,k) + (1-\pi^{exit}) \int_0^{\bar{\xi}} CV_t(z,k,\xi) d\xi \tag{8}$$

where CV_t^{exit} and CV_t denote the continuation values of exiting and surviving firms, respectively. With probability π^{exit}, a firm is forced to exit after the production stage. Exiting firms have the liquidation value

$$CV_t^{exit}(z,k) = (1-\delta(1-\chi)) q_t k. \tag{9}$$

as they do not need to pay adjustment costs. The continuation value of a surviving firm is

$$CV_t(z,k,\xi) = \max \{CV_t^a(z,k,\xi), CV_t^n(z,k)\}, \tag{10}$$

which reflects that surviving firms can decide whether to adjust their capital stock (CV_t^a) or not (CV_t^n). The continuation value of not adjusting is:

$$CV_t^n(z,k) = \mathbb{E}_t \left[\Lambda_{t+1} V_{t+1}(z', (1-\delta(1-\chi))k)\right] - q_t \chi \delta k, \tag{11}$$

while the continuation value of a firm that adjusts its capital stock is:

$$CV_t^a(z,k,\xi) = \max_{k'} \mathbb{E}_t \left[\Lambda_{t+1} V_{t+1}(z',k')\right] - q_t \left(k' - (1-\delta)k\right) - AC(k,k',\xi). \tag{12}$$

Policy Functions. The labor decision in equation (8) is static and independent of the capital decision, the optimal labor input being $n_t^*(z,k) = \left(\frac{p_t \nu z k^\theta}{w_t}\right)^{\frac{1}{1-\nu}}$. Thus, earnings net of labor costs are $\pi_t(z,k) \equiv p_t z k^\theta (n_t^*)^\nu - w_t n_t^*$.

The optimal capital decision is computed as follows. First of all, the solution to the maximization problem in equation (12) is the policy function $k_t^a(z,k)$, which is independent of ξ. This policy function allows us to compute $CV_t^a(z,k,\xi)$. Since, $CV_t^a(z,k,\xi)$ depends on ξ linearly, we can formulate a cutoff rule for the maximization problem in equation (10). Firms choose to adjust capital if and only if their fixed adjustment cost draw ξ is smaller or equal $\xi_t^T(z,k)$:

$$k_t^*(z,k,\xi) = \begin{cases} k_t^a(z,k) & \text{if } \xi \leq \xi_t^T(z,k) \\ (1-\delta(1-\chi))k & \text{if } \xi > \xi_t^T(z,k) \end{cases} \tag{13}$$

where $\xi_t^T(z,k) = \frac{CV_t^a(z,k,\xi=0) - CV_t^n(z,k)}{w_t}$. The hazard rate is defined as $\lambda_t(z,k) = \frac{\xi_t^T(z,k)}{\bar{\xi}}$ between 0 and 1.

3.2. New Keynesian Block

We separate nominal rigidities from the investment block of the model. A fixed mass of retailers $i \in [0,1]$ produces differentiated varieties $\widetilde{y}_{it}$ from the undifferentiated intermediate goods produced by the production firms. There is a one-to-one production technology $\widetilde{y}_{it} = y_{it}$, where y_{it} is the amount of the intermediate good that retailer i purchases. Retailers face Rotemberg quadratic price adjustment costs $\frac{\varphi}{2}\left(\frac{\widetilde{p}_{it}}{\widetilde{p}_{it-1}} - 1\right)^2 Y_t$, where $\widetilde{p}_{it}$ is the relative price of variety i. A representative final good producer aggregates the differentiated varieties optimally into the final good according to $Y_t = \left(\int \widetilde{y}_{it}^{\frac{\gamma-1}{\gamma}}\, di\right)^{\frac{\gamma}{\gamma-1}}$. The resulting demand function for retail good $\widetilde{y}_{it}$ is: $\widetilde{y}_{it} = \left(\frac{\widetilde{p}_{it}}{P_t}\right)^{-\gamma} Y_t$, where $P_t = \left(\int_0^1 \widetilde{p}_{it}^{1-\gamma} di\right)^{\frac{1}{1-\gamma}}$ is the price of the final good. We log-linearize the optimality condition of the retailer's problem to obtain the familiar New Keynesian Phillips Curve (NKPC):

$$\log(1+\pi_t) = \frac{\gamma-1}{\varphi} \log \frac{p_t}{p^*} + \beta \, \mathbb{E}_t \log(1+\pi_{t+1}) \tag{14}$$

where $\pi_t \equiv P_t/P_{t-1} - 1$ is the inflation rate, p_t is the relative price (in terms of the final good) of the intermediate good, and $p^* = \frac{\gamma-1}{\gamma}$ is its steady-state value.

3.3. Capital Good Producer

There is a representative capital good producer operating in a perfectly competitive market. It transforms units of the final good into new capital subject to external capital adjustment costs: $I_t = \left[\frac{\delta^{1/\kappa}}{1-1/\kappa}\left(\frac{I_t^Q}{K_t}\right)^{1-1/\kappa} - \frac{\delta}{\kappa-1}\right] K_t$, where I_t^Q represents the amount of the final good used, I_t the amount of new capital produced, and K_t is the total stock of capital in the beginning of period t. The parameter κ determines the strength of external capital adjustment costs. Optimal behavior implies that the relative price of capital (q_t) has to satisfy the following condition

$$q_t = \left(\frac{I_t^Q/K_t}{\delta}\right)^{1/\kappa} \tag{15}$$

3.4. Central Bank

The central bank sets the nominal interest rate r_t^n according to a Taylor rule

$$\log(1+r_t^n) = \rho_r \log(1+r_{t-1}^n) + (1-\rho_r)\left[\log\frac{1}{\beta} + \varphi_\pi \log(1+\pi_t)\right] + \epsilon_t^m \tag{16}$$

where ϵ_t^m is a monetary policy shock, ρ_r is the interest rate smoothing parameter, and φ_π is the reaction coefficient to inflation.

3.5. Household

There is a representative household, which consumes C_t^h, supplies labor N_t^h, and saves or borrows in one-period non-contingent bonds B_t^h. The household's objective is to maximize expected lifetime utility

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\log(C_t^h) - \psi N_t^h\right), \tag{17}$$

subject to the flow budget constraint $P_t C_t^h + Q_t^B B_t^h \leq B_{t-1}^h + W_t N_t^h + \Pi_t$, where Q_t^B is the nominal one-period risk-free bond price (one unit of B_t pays one unit of currency at $t+1$), W_t is the nominal wage, and Π_t subsumes additional transfers to and from the household.[16]

[16] Π_t includes dividends from intermediate good producers, retailers, and the final good producer, as well as the initial capital endowment k_0, which entering firms receive from the household.

Solving the household's optimization problem leads to the following optimality conditions

$$\Lambda_{t+1} = \beta \mathbb{E}_t \left[\frac{C_t^h}{C_{t+1}^h} \right] \tag{18}$$

$$w_t = \psi C_t^h \tag{19}$$

where Λ_{t+1} is the household's stochastic discount factor between periods t and $t+1$, and w_t is the real wage. Appendix E.1 defines an equilibrium in this economy.

4. Quantitative Results

We use the model to inspect the investment channel of monetary policy. After calibrating the model (Section 4.1), we analyze the effects of an expansionary monetary policy shock (Section 4.2), mirroring the empirical analysis. Thereafter, we study the aggregate implications of firm heterogeneity in the model (Section 4.3) and discusses how the empirical and model-based analyses enhance our understanding of the investment channel of monetary policy (Section 4.4).

4.1. Calibration

We calibrate the model to the U.S. economy. Wherever possible, we rely on data sources that are representative of the entire economy. We begin by fixing a subset of parameters to conventional values. These parameters are summarized in Table A.1. Given these fixed parameters, we fit the remaining parameters to match the moments listed in panel (b) of Table 1. The fitted parameters are listed in panel (a) of Table 1.

Since a model period corresponds to a quarter, the discount factor is set to $\beta = 0.99$. The labor disutility parameter is set to $\psi = 0.58$.[17] Capital and labor coefficients are set to standard values, that is, $\theta = 0.21$ and $\nu = 0.64$ (Ottonello and Winberry, 2020). The depreciation rate $\delta = 1.93\%$ generates an annual aggregate investment rate of 7.7% as reported in Zwick and Mahon (2017). We target the standard deviation of idiosyncratic TFP shocks σ_z, but fix their persistence ρ_z due to the identification problem discussed in Clementi and Palazzo (2015). We set ρ_z to 0.95 (Khan and Thomas, 2008). The exit probability π^{exit}

[17] This value follows from normalizing the steady-state real wage w to 1.

is set to 1.625% as in Koby and Wolf (2020).[18] We choose standard values for the parameters of the New Keynesian block, i.e. $\varphi = 90$ and $\gamma = 10$ (Ottonello and Winberry, 2020). The coefficient on inflation in the Taylor rule φ_π is set to 1.5 and the interest rate smoothing parameter ρ_r is set to 0.75. External capital adjustment costs κ are set to 11 to roughly match the peak effect of a monetary policy shock on investment relative to the peak effect on output documented empirically (Figure D.1).

The five parameters listed in panel (a) of Table 1 are chosen to match the five targeted moments listed in panel (b) of Table 1. Even though all parameters are calibrated jointly, we briefly explain which moments are particularly informative about which parameters. First, we target the standard deviation of investment rates, because it is informative about the volatility of idiosyncratic TFP shocks. Second, we target the average investment rate as it is informative about both adjustment cost parameters. Increasing either adjustment cost dampens investment rates in particular of young firms and therefore the average investment rate. Third, we target the autocorrelation of investment rates, because it is informative about the relative importance of fixed and convex adjustment costs. Convex adjustment costs generate a positive autocorrelation, whereas fixed adjustment costs generate a negative or zero autocorrelation. For these three moments, we use the statistics reported in Zwick and Mahon (2017). Fourth, we target the relative size of entrants, which is informative about the initial capital of entrants. This moment is computed from Business Dynamics Statistics (BDS) data. Fifth, we target the spike rate of old firms relative to the spike rate of young firms, which is informative about the maintenance investment parameter. The more depreciation is undone by maintenance investment, the less frequently do old firms need to make an extensive margin investment. Thus, a higher maintenance parameter leads to a lower spike rate among old firms. This moment needs to be computed from Compustat data, since it is the only data source which includes both investment rates and firm age.

[18]This exit probability brings the age distribution as close to the data as possible without using age-specific exit probabilities.

Table 1: Model Calibration

Parameter	**Description**	**Value**
σ_z	Volatility of TFP Shock	0.07
k_0	Initial Capital of Entrants	2.27
$\bar{\xi}$	Upper Bound on Fixed Adjustment Cost	0.90
ϕ	Convex Adjustment Cost	2.20
χ	Maintenance Investment Parameter	0.34

(a) Fitted Parameters

Moment	**Data**	**Model**
Standard Deviation of Investment Rates	0.20	0.18
Average Investment Rate	0.12	0.13
Autocorrelation of Investment Rates	0.38	0.38
Relative Size of Entrants	0.29	0.29
Relative Spike Rate of Old Firms	0.40	0.40

(b) Empirical & Simulated Moments

Notes: Data moments related to investment rates are taken from Zwick and Mahon (2017) (Appendix, Table B.1, Unbalanced Sample). The relative spike rate of old firms is computed from Compustat data. Corresponding model moments are computed from a simulation of a large panel of firms. The relative size of entrants is taken from Business Dynamics Statistics (BDS). In the model, this moment can be computed from the steady-state distribution.

Figure 5: Life-Cycle Profiles

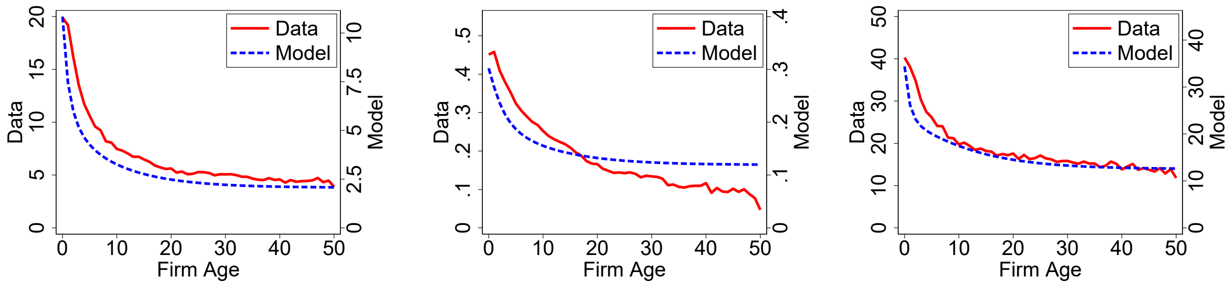

(a) Average Investment Rate (b) Frequency of Adjustment (c) Conditional Inv. Rate

Notes: All investment, hazard, and spike rates refer to a quarter and are averages across all firms of a given age. To compute the frequency of adjustment reported in panel (b), we use the spike rate (defined as the fraction of firms choosing an investment rate larger than 10%) in the data and the hazard rate in the model (defined as the fraction of firms choosing to pay the fixed adjustment cost and adjust their capital stock). Panel (c) plots the investment rate among all firms that make an adjustment as defined in panel (b).

4.1.1. Untargeted Moments

At the calibrated parameters, the simulated moments match the targeted empirical moments well (panel (b) of Table 1). Before moving to the main analysis, we highlight that the model is also capable of reproducing well-known facts regarding (i) firm life cycles, (ii) the aggregate effects of monetary policy shocks (Appendix E.2.1), and (iii) the interest-rate-elasticity of aggregate investment (Appendix E.2.2). The model being able to match these untargeted moments serves as an external validation for the calibration of the model.

Firm Life-Cycle Profiles. Figure 5 shows that the model matches the life-cycle profiles of firm investment behavior very well. Panel (a) shows that in the data as in the model, the average investment rate is highest for newborn firms and falls with age. Panels (b) and (c) decompose the pattern of the average investment rate into the frequency of adjustment (extensive margin) and the investment rate conditional on adjusting capital (intensive margin). Evidently, the observation that young firms have higher average investment rates is driven in part by a higher frequency of adjustment and in part by higher conditional investment

Figure 6: Effect of Monetary Policy on the Distribution of Investment Rates

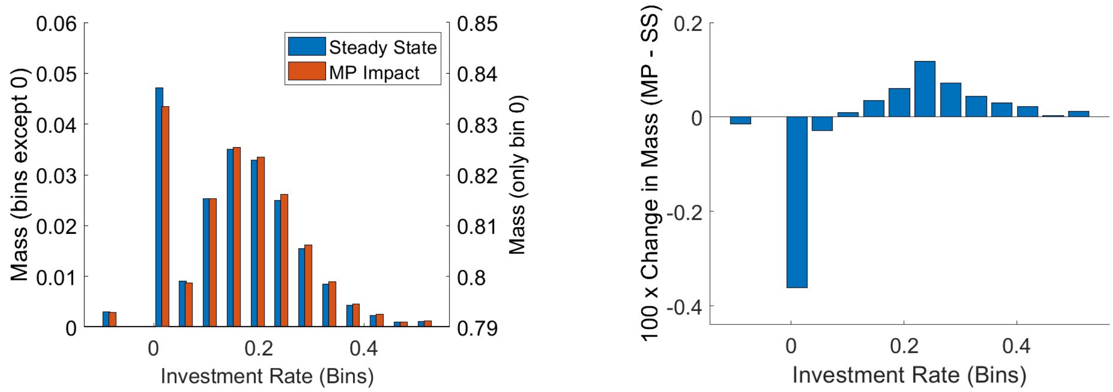

(a) Investment Rate Distributions

(b) Change in Distribution

Notes: Panel (a) plots the distribution of investment rates in steady state (blue bars) and after an expansionary monetary policy shock (red bars). To improve readability, the shock is scaled to reduce the nominal interest rate by 100 basis points. Panel (b) plots the difference between the two distributions shown in panel (a).

rates.[19] While the life-cycle profiles in the data and in the model align well, there are some differences in levels, which are discussed in Appendix E.3.

It is worth emphasizing that all three investment frictions are necessary to generate these life-cycle profiles. First, fixed adjustment costs generate lumpy investment behavior, so frequencies of adjustment below one, as shown in Panel (b). Second, convex adjustment costs ensure that young firms in the model choose plausible conditional investment rates, as shown in Panel (c), and do not immediately jump to their optimal size. Third, maintenance investment makes hazard rates decrease with age, as shown in Panel (b).

4.2. Monetary Policy and the Distribution of Investment Rates

Turning to the main analysis of this paper, Figure 6 plots the effect of a monetary policy shock on the distribution of investment rates.[20] Panel (a) plots the distribution of investment rates in steady state (blue bars) and in the period when an expansionary monetary policy shock has hit the economy (red bars), while panel (b) plots the difference between the two distributions. Evidently, monetary policy affects the distribution of investment rates: after an interest rate cut, fewer firms choose a small investment rate and more firms choose to make a sizeable investment. This observation corresponds to **Fact 1** documented in Section 2. Monetary policy affects the distribution of investment rates because it alters some firms' extensive margin investment decision. This is confirmed by panel (a) of Figure 7 which shows that the hazard rate rises and the inaction rate falls, matching the empirical evidence. In addition, the dispersion of investment rates increases, because the upper quantiles of the investment rate distribution respond much more than the lower quantiles (Figure A.9).

Monetary policy affects the average investment rate not only via the extensive margin investment decision but also via the intensive margin investment decision. To assess the relative importance of both margins, we decompose the effect on the average investment rate into contributions of the extensive and intensive margin, similar to the empirical exercise presented in Figure 4.[21] Panel (b) of Figure 7 shows that the expansionary monetary policy shock increases the average investment rate and also that the model attributes a significant portion of the change in the average investment rate to the extensive margin, as in the data.[22]

[19] In the model, we use the hazard rate to identify extensive margin adjustments unambiguously. In the data, we use the spike rate as an empirical proxy for the (unobservable) hazard rate. Using the spike rate instead of the hazard rate in the model, we obtain very similar life-cycle profiles and results (Figure A.12).

[20] The size of the monetary shock is chosen to roughly match the peak effects on output and investment seen in the data. This implies that the nominal interest rate falls by around 25 basis points on impact.

[21] This decomposition is computed by holding either hazard rates at steady-state levels (intensive margin) or conditional investment rates at steady-state levels (extensive margin), see equation (3).

[22] In Appendix E.4, we study the response of the average investment rate as well as the contributions of the extensive and intensive margin to two additional macroeconomic shocks: an aggregate TFP shock, and a wage-markup shock. This exercise shows that the contribution of the extensive margin can differ substantially if a shock affects the adjustment costs directly.

Figure 7: Effect of Monetary Policy on Hazard, Inaction & Average Investment Rate

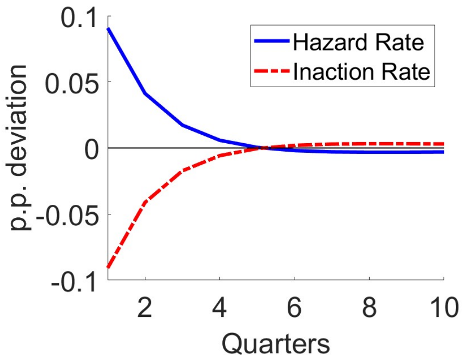

(a) Hazard & Inaction Rate

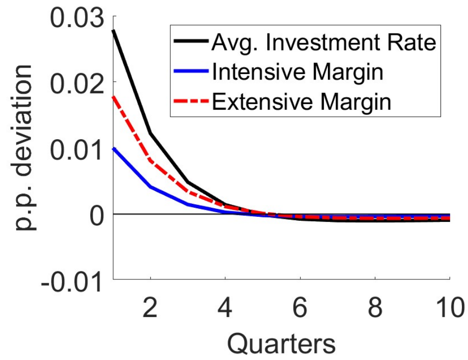

(b) Avg. Investment Rate

Notes: Panel (a) of this figure plots the effect of a monetary policy shock on the hazard and inaction rate of all firms. Panel (b) plots the effect of a monetary policy shock on the average investment rate of all firms and decomposes this effect into an extensive margin contribution and an intensive margin contribution.

Heterogeneous Effects: Young vs. Old Firms. In addition, the model reproduces the empirical finding that the effect of monetary policy on the distribution of investment rates is heterogeneous across age groups, as shown in Figure 8. Panels (a) and (b) plot the distribution of investment rates before and after an expansionary monetary policy shock of young and old firms, respectively. Panels (c) and (d) plot the changes in the distributions, highlighting that after an interest rate cut, the decrease in zero and small investment rates and the increase in large investment rates is more pronounced among young firms than among old firms. This corresponds to **Fact 2**. The heterogeneous effect on the distribution reflects that more young firms switch from being inactive to making a large investment. In line with this, panel (c) of Figure 9 shows that the hazard rate rises more strongly for young firms, confirming the stronger investment channel of monetary policy along the extensive margin.

In line with the heterogeneous effect on the distribution of investment rates, monetary policy affects *average* investment rates differently across age groups. Panel (a) of Figure 9 shows that after an expansionary shock, the average investment rate rises more strongly among young firms than among old firms. Panel (b) decomposes this heterogeneous effect into extensive and intensive margin contributions, similar to the empirical exercise shown in panel (b) of Figure 4. Evidently, the total difference is predominantly driven by the extensive

Figure 8: Effect of Monetary Policy on the Distribution of Inv. Rates (by Age Group)

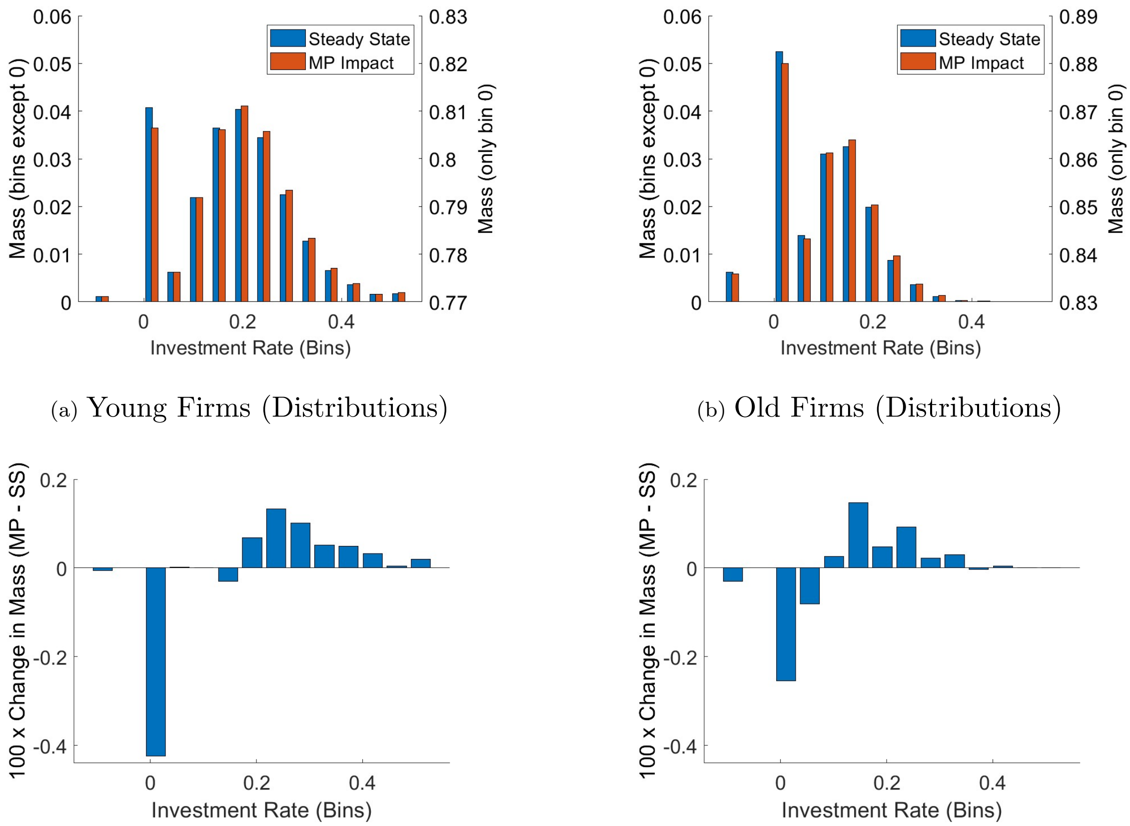

(a) Young Firms (Distributions)

(b) Old Firms (Distributions)

(c) Young Firms (Difference)

(d) Old Firms (Difference)

Notes: Panels (a) and (b) of this figure plot the distribution of investment rates of young (old) firms in steady state (blue bars) and after a monetary policy shock (red bars). To improve readability, the shock is scaled to reduce the nominal interest rate by 100 basis points. Panels (c) and (d) plot the difference of the two distributions for young (old) firms.

margin, corresponding to **Fact 3**. However, the *heterogeneous hazard rate increase*, shown in panel (c), is not the only reason for the heterogeneous effect along the extensive margin. The additional young adjusters also choose on average a higher investment rate than the additional old adjusters, which is also visible in panels (c) and (d) of Figure 8. We refer to this as the *heterogeneous size effect*. Panel (d) shows that both effects are quantitatively important.

Figure 9: Heterogeneous Effect of Monetary Policy (by Age Group)

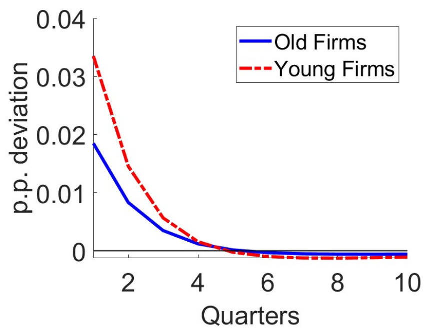

(a) Avg. Investment Rate

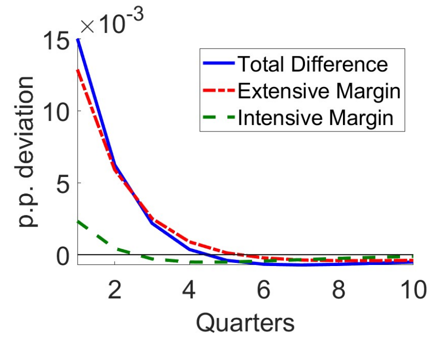

(b) Extensive vs. Intensive Margin

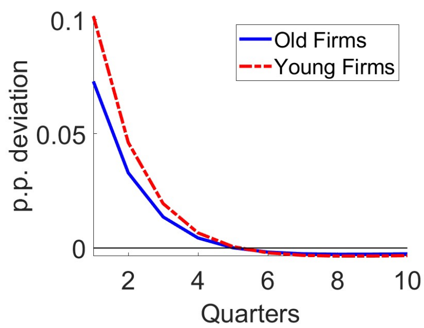

(c) Hazard Rate

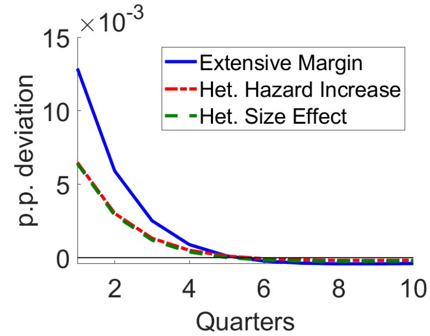

(d) Extensive Margin Decomposition

Notes: Panel (a) of this figure plots the effect of a monetary policy shock on the average investment rates of young and old firms. Panel (b) decomposes the differences of the two IRFs in panel (a) into an extensive margin contribution and an intensive margin contribution. Panel (c) plots the effect of a monetary policy shock on the hazard rates of young and old firms. Panel (d) further decomposes the IRF of the extensive margin contribution in panel (b) into the heterogeneous hazard rate increase and the heterogeneous size effect.

Intuition. We now build intuition for the extensive margin investment channel, while in Appendix B, we derive the *heterogeneous size effect* and the *heterogeneous hazard rate increase* analytically in a stylized two-period model. Drawing on the decomposition of the group-specific average investment rate already used in Section 2.5, the average investment rate ($\bar{i}$) among firms of group g is

$$\overline{i_g} = \psi_g i_g^* + (1 - \psi_g) i^m \tag{20}$$

where ψ_g is the hazard rate, i_g^* is the average investment rate among firms paying the fixed cost, and i^m is the time-invariant (maintenance) investment rate of firms not paying the fixed cost.

The interest rate sensitivity of the average investment rate is:

$$\frac{\partial \overline{i_g}}{\partial r} = \underbrace{\frac{\partial \psi_g}{\partial r}\left(i_g^* - i_g^m\right)}_{\text{Extensive Margin}} + \underbrace{\psi_g \frac{i_g^*}{\partial r}}_{\text{Intensive Margin}} \tag{21}$$

A decrease in the interest rate leads to a higher hazard rate ($\frac{\partial \psi_g}{\partial r} < 0$) because it increases the discounted benefit of investing while leaving the cost of investing (in partial equilibrium) unchanged, leading to more firms paying the fixed cost. This is what we label the extensive margin investment channel.

Comparing the increase in the average investment rate due to the extensive margin among young (Y) and old (O) firms, we uncover the two effects plotted in Panel (d) of Figure 9:

$$\begin{aligned} HetExt_{Y-O} &= \underbrace{\frac{\partial \psi_Y}{\partial r}\left(i_Y^* - i^m\right)}_{\text{Young Firms}} - \underbrace{\frac{\partial \psi_O}{\partial r}\left(i_O^* - i^m\right)}_{\text{Old Firms}} \\ &= \underbrace{\frac{\partial \psi_O}{\partial r}\left(i_Y^* - i_O^*\right)}_{\text{Heterogeneous Size Effect}} + \underbrace{\left(\frac{\partial \psi_Y}{\partial r} - \frac{\partial \psi_O}{\partial r}\right)\left(i_Y^* - i^m\right)}_{\text{Heterogeneous Hazard Rate Increase}} \end{aligned} \tag{22}$$

On the one hand, there is the *heterogeneous size effect*. Among the *new* adjusters, young firms choose higher investment rates conditional on adjusting than old firms ($i_Y^* - i_O^* > 0$). Panel (c) of Figure 5 has shown that in the data and in the model, young firms have on average higher conditional investment rates. Therefore, there would be a heterogeneous effect on average investment rates even if an interest rate cut had the same effect on hazard rates of young and old firms.

On the other hand, there is a *heterogeneous hazard rate increase* as an interest rate cut raises the hazard rate of young firms by more than the hazard rate of old firms. In general, hazard rates rise because the discounted benefit of investing rises, while the cost remains unchanged.[23] This increase in the discounted benefit of investing is larger for young firms.

[23] The discounted benefit of investing is $\frac{1}{1+r}\Big(V\big(k_0 \times (1+i^*)\big) - V\big(k_0 \times (1-\delta(1-\chi))\big)\Big)$. Due to general equilibrium effects, the cost of investing is affected by interest rate changes as well, but there is no direct (partial equilibrium) effect.

The reason is that young firms have a higher marginal product of capital, which reflects that young firms are farther away from their "optimal size" as they are on average smaller and the model features decreasing returns to scale. Therefore, young firms are more inclined to make an investment, and monetary policy can more easily induce them to do so.

The Role of Fixed and Convex Adjustment Costs. To generate a heterogeneous effect of monetary policy across age groups in the model, the fixed cost is the dominant adjustment cost. In isolation, the fixed adjustment cost generates 55% of the heterogeneous average effect (panel (a) of Figure 9), while the convex adjustment cost only generates 29%. This nicely illustrates that also the interaction of both adjustment costs is important, accounting for the remaining 16%. The heterogeneous size effect, derived in equation (22), helps to understand why this interaction matters. This effect requires a change in hazard rates and therefore a fixed adjustment cost, but also different conditional investment rates, which depend on the convex adjustment cost.

4.3. Aggregate Implications

The features of firm-level investment behavior—lumpiness and life-cycle dynamics—that we emphasize to be important to match the empirical facts have aggregate implications. Lumpy investment behavior implies that monetary policy is particularly effective in stimulating investment whenever there are many firms that are "close to paying the fixed cost", as they can be induced to make a meaningful investment. As also age matters for investment behavior, our model implies that the effectiveness of monetary policy varies over the business cycle as well as due to long-run trends in the age distribution of firms.

First, firm aging affects the investment channel of monetary policy. Specifically, monetary policy is more effective when business dynamism is high (high entry and exit rates). This is because when business dynamism is high, the share of young firms in the economy is higher, and young firms are more easily induced to pay the fixed cost and invest. In our baseline calibration, the share of entrants (firms of age 0) is around 6.5% and thus close to the value observed in the U.S. over the past decade. To quantify the relevance of the decline in business dynamism, documented by Haltiwanger et al. (2012) among others, for monetary policy, we compute a counterfactual "high-dynamism" calibration, which features a twice as high share of entrants, as observed in 1984 (13%).[24] We also compute a "low-dynamism" calibration, which features a 50% lower share of entrants (3.375%). In the steady state of the high-dynamism (low-dynamism) calibration, there are relatively more (fewer) young and therefore small firms, as panel (a) of Figure A.11 shows. Since young firms are more responsive to interest rate changes, the impact effect of a monetary policy shock on aggregate investment, which is 1.44% in the baseline calibration, is with 1.61% around 11.5% larger in the high-dynamism calibration and with 1.32% around 8.5% smaller in the low-dynamism calibration. Panel (b) of Figure A.11 shows that these differential effects are persistent. Hence, according to our model, the well-documented decline in business dynamism has made monetary policy less effective in stimulating investment.

Second, monetary policy is more effective in a boom than in a bust. In a boom, productivity and the return on capital are high and therefore many firms are "close to paying the fixed cost" and can be induced to invest. Figure A.10 shows that the increase in investment following the same monetary policy shock is about 22% more effective in a large boom than in a deep recession. This aligns well with the findings of Koby and Wolf (2020).

In a general equilibrium framework, the two aggregate implications are not obvious from the empirical results, making a model-based quantification important. In partial equilibrium, a higher share of young firms (or, a boom) naturally implies a larger increase in investment demand in response to an expansionary monetary policy shock. In general equilibrium, however, the larger increase in investment demand raises the price of capital, which in turn crowds out investment demand to some extent. Hence, a higher share of young firms does not necessarily imply that monetary policy is more effective, even when there are heterogeneous effects across age groups. As Koby and Wolf (2020) explain in detail, state-dependent effects only arise if the price elasticity of aggregate investment, which determines the extent of the crowding-out, is sufficiently low, as is the case in our model (see Appendix E.2.2).

[24]To do so, we double both entry and exit rates while holding all other model parameters fixed. This ensures that the mass of firms remains 1, but the distribution of firms features more young firms.

4.4. Implications for Theories of the Investment Channel of Monetary Policy

Plenty of empirical evidence documents that expansionary monetary policy stimulates aggregate investment. However, it is much less clear which groups of firms drive this increase in investment in response to monetary policy and why. A growing literature aims to gain a better understanding of the investment channel by exploiting firm-level data to tease out which mechanisms make firms particularly responsive and which frictions limit firms' responsiveness to monetary policy (Ottonello and Winberry, 2020; Jeenas, 2023; Cloyne et al., 2023). The three empirical facts presented in this paper, paired with insights from the heterogeneous-firm model with lumpy investment, offer several implications that are helpful for the understanding of the investment channel of monetary policy.

Fixed Adjustment Costs. The change in the distribution of investment rates after a monetary policy shock (Fact 1, documented in Figure 1) highlights the importance of fixed adjustment costs and the extensive margin investment decision for understanding the investment channel. As discussed in detail in Appendix E.5, a model without fixed adjustment costs does not generate a change in the distribution of investment rates that is in line with the evidence. Figure E.4 shows that without fixed adjustment costs, the change in the distribution displays neither an outsized decrease in the share of firms making a very small investment, nor a sizeable increase in the share of very large investments, nor the absence of meaningful changes among negative investment rates.

Financial Acceleration. The differential change in the distribution of investment rates among young and old firms after a monetary policy shock (Fact 2, documented in Figure 2) offers insights about financial frictions and financial acceleration. This is because young firm age is oftentimes used as a proxy variable for tight financial constraints and a heterogeneous effect of monetary policy across age groups is interpreted as evidence for financial acceleration (Cloyne et al., 2023). Our model shows that financial acceleration is not necessary to generate a heterogeneous effect across age groups. Hence, there is an issue of observational equivalence as we show that a model with a non-financial friction (fixed adjustment costs) equally generates a heterogeneous effect across age groups.[25]

Moreover, our evidence shows that around 60% of the heterogeneous effect of monetary policy across age groups is explained by the extensive margin (Fact 3, documented in Figure 4). This evidence cannot be rationalized with the classical financial accelerator in the spirit of Bernanke et al. (1999), which operates along the intensive margin.

5. Conclusion

In this paper, we document that monetary policy reshapes the distribution of investment rates. An interest rate cut leads to fewer zero and small investment rates and more large investment rates, especially among young firms. This evidence highlights two features of firm-level investment behavior that are important to understand the investment channel of monetary policy. First, firm-level investment is lumpy and therefore, there is a quantitatively relevant investment channel along the extensive margin. Second, life-cycle investment dynamics are important to understand the heterogeneous effects of monetary policy on firms of different age groups. Young firms tend to grow whereas many old firms have reached their desired size. Therefore, young firms can more easily be induced to invest, leading to a higher responsiveness to monetary policy. We build a heterogeneous-firm model with fixed adjustment costs and firm dynamics not only to rationalize these empirical findings but also to investigate the aggregate implications of the key features of the model—lumpy investment and firm life cycles. We show that the effectiveness of monetary policy varies over the business cycle as well as due to long-run trends. In particular, the secular decline in business dynamism has weakened the investment channel.

Our work highlights an important avenue for future research, guided by three questions: Why do young firms grow slowly? How are the relevant frictions affected by economic policy? In turn, how are the effects of economic policies determined by these frictions? A long literature has emphasized financial frictions as the key constraint for young firms (Gertler and Gilchrist, 1994; Cloyne et al., 2023). We show that fixed capital adjustment costs—generating

[25]This argument is reminiscent of Gomes (2001) who shows that financial frictions are neither necessary nor sufficient to estimate a cash-flow effect in reduced-form regressions of investment.

lumpy investment behavior—are another key constraint for young firms and determine the effectiveness of monetary policy. Yet, there are further non-financial factors that constrain particularly young firms, such as uncertainty about productivity and demand (Jovanovic, 1982; Chen et al., 2023). Gaining a better and more complete understanding of why young firms grow slowly and how the relevant frictions matter for economic policy is crucial to guiding the design of effective policy interventions in the future.

References

Bachmann, R., Bayer, C., 2014. Investment Dispersion and the Business Cycle. American Economic Review 104, 1392–1416.

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, 29–67.

Baley, I., Blanco, A., 2021. Aggregate dynamics in lumpy economies. Econometrica 89, 1235–1264.

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

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

Caballero, R.J., Engel, E.M., Haltiwanger, J.C., 1995. Plant-level adjustment and aggregate investment dynamics. Brookings Papers on Economic Activity , 1–54.

Chen, C., Senga, T., Sun, C., Zhang, H., 2023. Uncertainty, imperfect information, and expectation formation over the firm's life cycle. Journal of Monetary Economics .

Christiano, L.J., Eichenbaum, M., Evans, C.L., 2005. Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. Journal of Political Economy 113.

Clementi, G.L., Palazzo, B., 2015. On The Calibration of Competitive Industry Dynamics Models. Mimeo, New York University .

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

Cooper, R.W., Haltiwanger, J., 2006. On the Nature of Capital Adjustment Costs. Review of Economic Studies 73, 611–633.

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

Fang, M., 2023. Lumpy Investment, Fluctuations in Volatility and Monetary Policy. Mimeo, University of Florida .

Farre-Mensa, J., Ljungqvist, A., 2016. Do measures of financial constraints measure financial constraints? Review of Financial Studies 29, 271–308.

Gertler, M., Gilchrist, S., 1994. Monetary policy, business cycles, and the behavior of small manufacturing firms. Quarterly Journal of Economics 109.

Gertler, M., Karadi, P., 2015. Monetary Policy Surprises, Credit Costs and Economic Activity. American Economic Journal: Macroeconomics 7, 44–76.

Gomes, J.F., 2001. Financing investment. American Economic Review 91, 1263–1285.

González, B., Nuño, G., Thaler, D., Albrizio, S., 2022. Firm heterogeneity, capital misallocation and optimal monetary policy. Banco de Espana Working Paper 2145 .

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

Haltiwanger, J., Jarmin, R., Miranda, J., 2012. Where have all the young firms gone? Business Dynamics Statistics Briefing 6.

Jeenas, P., 2023. Firm Balance Sheet Liquidity, Monetary Policy Shocks, and Investment Dynamics. Mimeo, Universitat Pompeu Fabra .

Jovanovic, B., 1982. Selection and the evolution of industry. Econometrica .

Jungherr, J., Meier, M., Reinelt, T., Schott, I., 2022. Corporate debt maturity matters for monetary policy. Mimeo, University of Bonn .

Khan, A., Thomas, J., 2008. Idiosyncratic shocks and the role of nonconvexities. Econometrica 76, 395–436.

Khan, A., Thomas, J.K., 2003. Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter? Journal of Monetary Economics 50, 331–360.

Koby, Y., Wolf, C., 2020. Aggregation in heterogeneous-firm models: Theory and measurement. Mimeo, Princeton University .

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

Loria, F., Matthes, C., Zhang, D., 2023. Assessing macroeconomic tail risk. Available at SSRN 4002665 .

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

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

Reiter, M., Sveen, T., Weinke, L., 2020. Agency costs and the monetary transmission mechanism. B.E. Journal of Macroeconomics 20, 1–11.

Tenreyro, S., Thwaites, G., 2016. Pushing on a string: Us monetary policy is less powerful in recessions. American Economic Journal: Macroeconomics 8, 43–74.

Thomas, J.K., 2002. Is lumpy investment relevant for the business cycle? Journal of Political Economy 110, 508–534.

Winberry, T., 2021. Lumpy investment, business cycles, and stimulus policy. American Economic Review 111, 364–96.

Zwick, E., Mahon, J., 2017. Tax Policy and Heterogeneous Investment Behavior. American Economic Review 59, 379–388.