The Investment Channel of Monetary Policy: Disentangling Firm Heterogeneity*

Thomas Drechsel
Johns Hopkins University, NBER, CEPR

Daniel Lewis
University College London, CEPR

Davide Melcangi
Federal Reserve Bank of New York

Laura Pilossoph
Duke University, NBER

September 11, 2026

Abstract

To study monetary policy transmission at the firm level, researchers typically consider heterogeneity along a small number of firm characteristics, such as firm size or leverage. We instead estimate the full distribution of firms' investment responses to changes in monetary policy, using a clustering regression framework. Our novel approach can capture multidimensional and unobservable heterogeneity across firms and time. We find that investment by most firms in most time periods responds little to monetary policy. Only about 5% of firm-time observations are associated with a strong response. In those cases, a 25 basis point interest rate hike lowers the quarterly growth rate of firm capital by about 1 percentage point on average. We then correlate our investment sensitivity estimates with observable firm characteristics. While typical characteristics studied in the literature predict sensitivity, we also uncover novel correlates. In general, the responsiveness of investment cannot easily be explained with a small number of standard firm-level observables. We demonstrate that our empirical estimates provide new and critical targets for disciplining structural models of firm heterogeneity and monetary policy.

Keywords: Monetary policy, investment, firm heterogeneity, panel local projections, clustering. **JEL Classification:** D25, E22, E32, E52.

*The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. We thank Tincho Almuzara, Adrien Auclert, Mark Bils, Luigi Bocola, Keshav Dogra, Simon Gilchrist, Kilian Huber, Silvia Miranda-Agrippino, Matthias Meier, Manuel Menkhoff (discussant), Pablo Ottonello, Gianluca Violante, Thomas Winberry, and Chris Wolf for detailed discussions. Conference and seminar participants at the NBER Summer Institute, Ohio State University, University of Manchester, the New York Fed, the Richmond Fed, De Nederlandsche Bank, the SED in Copenhagen, CEPR ESSIM Lille, the Research Conference in Honor of Raf Wouters, the IMAEF conference in Syros, the annual CEBRA meeting in Copenhagen, the CEPR Conference on New Challenges in Monetary Economics, and the 4th Workshop on Firm Heterogeneity and Macroeconomics in Bonn provided helpful comments. Augustin Belin, Miles Guerin, and Srinidhi Narayanan provided excellent research assistance.

1 Introduction

A central channel through which monetary policy affects economic activity is by altering firms' investment incentives. Recent research on the investment channel of monetary policy focuses on understanding heterogeneity in the cross section of firms. Typically, researchers choose a firm-level characteristic of interest, such as firm size or leverage, and study whether firm-level investment responds more or less strongly to a monetary policy surprise as this characteristic varies. This approach has generated rich insights into how firms' investment responses to monetary policy vary with firm size (Gertler and Gilchrist, 1994), leverage and distance to default (Ottonello and Winberry, 2020), liquidity (Jeenas, 2023), firm age (Cloyne et al., 2023), types of collateral (Caglio et al., 2024), or debt maturity (Jungherr et al., 2024). It is conceivable, however, that all of these firm characteristics explain the responsiveness of firms' investment to monetary policy. Moreover, these characteristics might interact, and potentially many additional and possibly unobservable firm characteristics might matter.

The goal of this paper is to uncover firm-level heterogeneity in the investment channel of monetary policy in a more comprehensive manner. We accomplish this by taking a novel approach relative to the literature. We remain agnostic *ex ante* about which firm characteristics are important for predicting the firm-level sensitivity of investment to monetary policy, and instead employ a statistical approach to characterize the full distribution of these sensitivities. Shifting the focus to the full distribution rather than the partial derivative across one or a few pre-selected firm features enables us to study heterogeneity that might be highly multidimensional, nonlinear, or unobservable. Moreover, our approach can characterize firm-level heterogeneity that varies across different firms or for a given firm over time.

The estimated heterogeneity we uncover can be linked to observable firm characteristics *ex post*. Such an analysis allows us to investigate the degree to which firm characteristics studied in the literature contribute to overall heterogeneity. We study the relevance of typical firm observables computed from firms' balance sheets and financial statements, such as size and leverage. We also consider firm characteristics from novel sources, such as firms' earnings calls or CFO surveys, which have so far received less attention in the literature on the investment channel of monetary policy, and which also matter for how firms respond to interest rate changes.

Our new approach for disentangling firm heterogeneity is a clustering (Gaussian mixture) regression framework, applied to the panel local projection setup commonly used in the existing literature. Our econometric method assigns firm-time observations from a panel dataset to different groups. Each group corresponds to one level of *responsiveness of*

investment to monetary policy, which for brevity we refer to as the "RIMP". Each group is associated with a probability (or frequency) that a firm-time observation belongs to it. The number of groups is finite and can be estimated using statistical selection criteria. The RIMP estimates and probabilities for each group, as well as their associated standard errors, allow us to provide a detailed characterization of heterogeneity in investment responses to monetary policy. Furthermore, using the group estimates and their probabilities, we can construct a posterior estimate for the predicted RIMP of each firm-time observation.

A similar clustering algorithm has been successfully applied to study household-level heterogeneity in the marginal propensity to consume (MPC) by Lewis et al. (2026). In this paper, we employ and further develop this approach to uncover firm-level heterogeneity in the RIMP. We apply the clustering methodology to US Compustat data and the panel local projection setup of Ottonello and Winberry (2020), a central paper in the literature on the investment channel of monetary policy. For comparability with the literature, most of our results focus on the impact response of investment to an estimated monetary policy surprise. Since longer horizon responses of investment are of particular interest in the broader macroeconomics literature, we also construct impulse response functions of investment in combination with our clustering procedure.

We present five sets of findings. The first is our estimated distribution of investment responses to monetary policy across firm-time observations, or the RIMP distribution. We find that most firm-time observations correspond to an economically small responsiveness. Notably, our estimates imply that even this low responsiveness is statistically distinct from zero, so all firms respond at least slightly to an interest rate change. About 5% of firm-time observations correspond to a large responsiveness, which is about 8 times larger than the smallest one. The magnitude of the largest RIMP implies that a 25 basis point (bp) interest rate hike lowers the quarterly growth rate of firm capital by about 1 percentage point (pp) on average. Compared to the literature, our estimate corresponding to the strongest responsiveness represents an economically large effect.

The second set of findings reveals the sources of variation in the RIMP. 20.6% of RIMP variation is explained by time-invariant differences across firms, 0.3% by aggregate movements common across firms, and 79.1% by firm-specific variation over time. In particular, the state of being associated with the strongest RIMP is highly transient. A firm that responds with a 1 pp investment cut to a 25 bp interest rate hike in a given quarter will respond with the same sensitivity in the next period with only 12% probability. A firm that responds with the smallest RIMP, on the other hand, will respond similarly in the following period with a 59% probability. In combination with the first set of findings, these patterns reveal that the investment of most firms in most time periods responds only weakly

to changes in monetary policy, while a few firms in some periods respond very strongly. We verify that this pattern is distinct from the unconditional investment behavior of firms, which previous research has shown to be lumpy.

The third set of findings examines the observable drivers of heterogeneity in the investment channel of monetary policy. We relate our RIMP estimates to firm-level characteristics *ex post*, by regressing the estimated RIMPs on firm-level observables. We begin this analysis by focusing on "traditional" firm features that have received attention in the literature, such as size, age, distance to default, or debt maturity. We show that these firm characteristics indeed play an important role in explaining heterogeneity. For instance, consistent with Gertler and Gilchrist (1994), smaller firms display a stronger investment sensitivity. Additionally, consistent with Cloyne et al. (2023), younger firms are more sensitive to monetary policy. In some cases, however, we find that the direction of the relationship differs from what the literature has found. We explain that such differences arise either because the local linear estimates in the literature miss important nonlinearities, or because there is underlying heterogeneity for which researchers cannot appropriately control.

Interestingly, a sizable share of variation in investment responsiveness remains unexplained by observable firm characteristics, even when jointly considering characteristics that have been proposed by the literature. A regression with the RIMP on the left-hand side and several observable firm characteristics on the right-hand side, as well as industry-time fixed effects, leads to a modest R^2 (around 5%). Even firm fixed effects or a LASSO regression that includes a large number of variables and allows for nonlinearities and interactions between them cannot increase the R^2 above 21%. Our list of variables includes firm observables that are motivated by the broader firm dynamics and investment literature, such as measured total factor productivity and indicators of investment inaction and investment spikes. We discuss the implications of this unobserved heterogeneity in the RIMP for the literature on monetary policy and firm investment.

In addition to traditional candidate drivers from the literature, in our fourth set of findings we turn to recent influential studies on corporate decision making. These studies go beyond the information about firms' financial statements and balance sheets that can be found in Compustat. For example, Gormsen and Huber (2025) construct measures of firms' perceived cost of capital and perceived discount rates from corporate earnings calls. They document a significant wedge between these two objects, which is strongly correlated with firm-level investment. We find that the RIMP is strongly negatively correlated with the behavioral wedges uncovered by that new line of research. We also document that firms with more volatile discount rates (i.e. less sticky discount rates) have a more negative RIMP, consistent with the model developed by Fukui et al. (2024).

Given that a variety of firm observables matter in the investment channel of monetary policy at the firm level, we investigate the multidimensionality of the RIMP drivers more explicitly. One way to illustrate the multidimensionality is by constructing bivariate RIMP distributions, which reveal a surface of heterogeneous RIMPs. For example, we show that small *and* young firms are very sensitive to monetary policy surprises. The same is true for small *and* less leveraged firms.

The fifth set of findings relates to the RIMP at longer horizons. These results demonstrate that our approach is applicable to a general panel local projection framework in which an impulse response function of investment to monetary policy shocks is estimated at the firm level. We also find considerable heterogeneity at longer horizons, with sizable responses persisting for some firms. Two years after a monetary policy surprise, the bulk of the RIMPs remain substantially negative. As with the impact responses, we project the predicted RIMP at different horizons on observable firm characteristics. The results are similar to the initial investment response, with small, young, and liquid firms being the most elastic. We also emphasize multidimensional RIMP heterogeneity at longer horizons.

Our results provide new and relevant empirical targets for structural models of heterogeneity in firm investment and monetary policy. To demonstrate their importance, we confront the structural model of Winberry et al. (2025) with our findings. We show that two versions of this model, which differ in the strength of capital adjustment costs and financial frictions, imply the same correlation between firm size and the RIMP. Hence, when informed by a univariate empirical relationship between firm size and the RIMP, both calibrations would be observationally equivalent. However, the two calibrations yield RIMP distributions with very different shapes. Since financial constraints and adjustment costs are key components of the monetary policy transmission mechanism, distinguishing heterogeneous firm models in these dimensions is crucial for the research agenda on monetary policy and investment. Only the RIMP distribution separates different configurations of these firm frictions in our model application, making it a relevant model target above and beyond existing targets put forward by the literature. While we focus on firm size, the conclusions from the model exercise also apply using correlations with other firm characteristics. In sum, our choice to estimate the full distribution of firms' investment responses to interest rate changes is not only relevant for understanding firm behavior in the transmission of monetary policy in its own right. Rather, it also provides critical empirical targets for researchers developing models of the investment channel of monetary policy.

Contributions to the literature. We contribute to several strands of research. First, our findings expand the literature that investigates the investment channel of monetary

policy empirically at the firm level. Ottonello and Winberry (2020) is a widely-cited paper that studies the role of financial frictions, in the form of variation in leverage and distance to default, in the transmission from monetary policy to investment. We follow their local projection setup in terms of the choice of data and specification. Earlier contributions along similar lines include Gertler and Gilchrist (1994), who focus on firm size. Jeenas (2023), Cloyne et al. (2023), Caglio et al. (2024), and Jungherr et al. (2024) provide additional insights. All of these papers select a firm-level characteristic *ex ante* and include it as an interaction term, while our approach uncovers the full RIMP distribution first, and projects the RIMP estimates on observable drivers ex post.

Second, methodologically we contribute to the body of work using clustering regression frameworks to understand economic heterogeneity, including unobserved heterogeneity.[1] Lewis et al. (2026) employ a similar approach to study household MPCs. The Gaussian mixture approach has its origin in Quandt (1972). Alternative approaches to clustering local projection coefficients include Bandeira (2023) (classifier LASSO) and Huang (2021) (K-means). We explain these alternative approaches in more detail in the main text. A key difference from most of them is that we allow the RIMP to be time-varying for a given firm. In our view, this is both a plausible and important assumption, as in many structural models the RIMP will depend on idiosyncratic state variables that vary over time. Methodologically, we also shed new light on how panel local projection approaches can be refined more generally. For a recent systematic review of these approaches, see Almuzara and Sancibrián (2024).

Third, the insights from our focus on the full RIMP distribution have implications for the broader firm dynamics and investment literature, in particular quantitative models of firm heterogeneity and monetary policy. Examples include, but are not limited to, Abel and Eberly (1996), Cooper and Haltiwanger (2006), Khan and Thomas (2013), Baley and Blanco (2021), and Winberry et al. (2025).

2 Methodology

This section introduces the firm-level local projection framework and explains how our clustering algorithm is applied to it. We also discuss various specification choices and practical considerations for our econometric methodology.

[1] An alternative way to go beyond observable heterogeneity would be to elicit hypothetical behavior in surveys, as recently done by Best et al. (2026). In the context of households, Colarieti et al. (2024) use clustering algorithms to classify hypothetical responses to income windfalls.

2.1 Empirical framework

We follow Ottonello and Winberry (2020) (henceforth OW) as closely as possible. We use the same Compustat Quarterly data from 1991:Q1 to 2007:Q4, and the same definitions of left-hand-side and right-hand side variables. Consider first the homogeneous panel local projection

$$\Delta^{+h} \log k_{i,t} = \alpha_i^h + \lambda^h \epsilon_t^m + \sum_{l=1}^{4} \Gamma^{h\prime} Z_{t-l} + \eta_{i,t}^h, \ h = 0, \ldots, H-1, \tag{1}$$

where i indexes firms, t indexes quarters, and h is the horizon of the local projection. We define the operator Δ^{+h} as the difference between time $t+h$ and $t-1$, so that $\Delta^{+h} \log k_{i,t} = \log k_{i,t+h} - \log k_{i,t-1}$. ϵ_t^m is a high-frequency monetary policy surprise, and Z_{t-l} is a vector of aggregate controls, lagged by l quarters. α_i^h is a firm fixed effect. λ^h is the parameter of interest. It determines the response of the change in firms' capital stock to a monetary policy surprise, that is, the RIMP for horizon h. $\eta_{i,t}^h$ is an iid normal error term.

Following OW, $k_{i,t}$ is measured as the book value of the tangible capital stock and the monetary policy surprise is computed as the change in the current-month Fed Funds Futures contract in a window that begins 15 minutes before and ends 45 minutes after an FOMC announcement, with an adjustment for the timing of the given meeting within the month. As in OW, the aggregate controls Z_{t-l} include real GDP growth, the unemployment rate, and CPI inflation.

Note that, differently from OW, Equation (1) does not include firm-level controls. As the goal of our econometric strategy is to uncover the full range of RIMP heterogeneity, we do *not* want to control for any firm-level characteristic *ex ante*. The same argument applies to sector-by-time fixed effects, which are included in the original OW specification but not included in ours. Instead, we will consider firm-level controls and sector-by-time fixed effects as possible determinants of the RIMP *ex post*. We only include a firm fixed effect α_i^h that absorbs a firm-specific mean and thus controls for any time-invariant differences in the growth rate of capital.

A key goal of the literature is to estimate heterogeneity in λ^h in Equation (1). For this purpose, a typical specification would be

$$\Delta^{+h} \log k_{i,t} = \alpha_i^h + \beta^h x_{i,t-1} \epsilon_t^m + \xi^h \epsilon_t^m + \delta^h x_{i,t-1} + \sum_{l=1}^{4} \Gamma^{h\prime} Z_{t-l} + \eta_{i,t}^h, \ h = 0, \ldots, H-1, \tag{2}$$

where $x_{i,t-1}$ is a (potentially time-varying) firm-level characteristic, such as firm size or distance to default. In this case, β^h captures the change in the response of investment to

a monetary policy surprise *as* $x_{i,t-1}$ *increases by one unit*. In other words, β^h captures heterogeneity in the RIMP, but only in the $x_{i,t-1}$-dimension. Different variations of this interaction term approach in Equation (2) can be found in the literature. For example, $x_{i,t-1}$ might be replaced by a dummy variable or might be expressed in deviations from a firm-specific mean, i.e. $(x_{i,t-1} - \mathbb{E}_j x_{i,s})$. Common to all those specifications, however, is that $x_{i,t-1}$ is a scalar or a low-dimensional vector, and chosen *ex ante* by the researcher.

2.2 The clustering approach: estimating RIMP heterogeneity

Instead of selecting a given $x_{i,t-1}$ *ex ante*, our approach seeks to characterize the entire unconditional RIMP distribution. We outline our clustering via GMLR approach by first explaining the general idea and then presenting the formal procedure and estimator, comparing it to alternatives where appropriate. For now, we focus on the impact response $h = 0$ and drop any reference to h to simplify the notation. We discuss the $h > 0$ case further below.

Main idea. In principle, we want to allow λ^h to be different for each (i,t) observation, so that the RIMP is firm-specific and time-varying, and then characterize the distribution of the estimated $\lambda^h_{i,t}$. However, identifying an (i,t)-specific slope coefficient from Equation (1) is infeasible, since the data only varies in the firm and time dimension. With an (i,t)-specific parameter, there would be one observation per parameter value, leaving no way to separate the value of the parameter from noise. Therefore, the econometrician must reduce the dimensionality of the problem. Our approach assigns observation (i,t) to groups $g = 1, ..., G$ with some probability π_g, where each group has a different λ_g and G is the total number of groups. G is taken as given for now; as we discuss below, it can be estimated using statistical criteria. Using the data, estimates of π_g and λ_g, and a Gaussianity assumption, it is then also possible to construct a firm-time-specific posterior prediction of the RIMP $\hat{\lambda}_{i,t}$. In what follows, we explain the formal procedure for implementing this approach. We then contrast it with alternative ways to achieve the desired dimensionality reduction. In robustness exercises, we also allow for group-specific intercepts, α_g.

Adjustments to the empirical framework. To estimate specification (1) with GMLR, we implement two pre-processing steps. First, we residualize the monetary policy surprises and other variables with respect to the aggregate controls, Z_{t-l}. This helps to ensure exogeneity of the monetary policy surprises (see, e.g., Bauer and Swanson, 2023). Relative to the Bauer and Swanson (2023) approach of orthogonalizing only the policy instrument, orthogonalizing all variables has no impact on the exogeneity condition, but improves

efficiency. Second, we absorb firm fixed effects, α_i^h, by demeaning both the left and the right-hand side of the estimation equation. We denote the resulting orthogonalized variables with a "$\widetilde{\cdot}$". The model we estimate is

$$\widetilde{\Delta \log k_{i,t}} = \sum_{g=1}^{G} d_{it,g} \lambda_g \widetilde{\epsilon_t^m} + \widetilde{\eta}_{i,t}, \quad \widetilde{\eta}_{i,t} \sim \mathcal{N}\left(0, \sigma_g^2\right), \tag{3}$$

where Δ is the one-period difference operator. $d_{it,g} = \mathbb{1}[it \in g]$ is a dummy variable that equals one if observation (i,t) belongs to group g.

Estimation procedure. After absorbing firm fixed effects, we make no further use of the panel structure when estimating the horizon-0 RIMP. Therefore our sample is essentially a cross section. Thus, for ease of notation, we denote an (i,t) observation simply as j. The complete-data likelihood is

$$L(\cdot) = \prod_{j=1}^{N} \prod_{g=1}^{G} \left(\pi_g \phi \left(\widetilde{\Delta \log k_j}; \lambda_g \widetilde{\epsilon_t^m}, \sigma_g^2 \right) \right)^{d_{jg}},$$

where π_g is the unconditional probability of belonging to $g \in G$ and $\phi(\cdot)$ is the normal probability density function. d_{jg} is a latent variable.

The incomplete data log-likelihood, based on the observable data, is

$$l(\cdot) \equiv \mathbb{E}_D \left[\log L(\cdot) \right] = \sum_{j=1}^{N} \sum_{g=1}^{G} \gamma_{jg} \left(\log \pi_g + \log \phi \left(\widetilde{\Delta \log k_j}; \lambda_g \widetilde{\epsilon_t^m}, \sigma_g^2 \right) \right),$$

where

$$\gamma_{jg} = \Pr\left(d_{jg} = 1 | \widetilde{\Delta \log k_j}, \widetilde{\epsilon_t^m} \right) = \frac{\pi_g \phi \left(\widetilde{\Delta \log k_j}; \lambda_g \widetilde{\epsilon_t^m}, \sigma_g^2 \right)}{\sum_{q=1}^{G} \pi_q \phi \left(\widetilde{\Delta \log k_j}; \lambda_q \widetilde{\epsilon_t^m}, \sigma_q^2 \right)}$$

is the posterior probability of firm-time observation j belonging to group g. This reflects the econometrician's uncertainty. We maximize this log-likelihood via the expectation-maximization (EM) algorithm. See Lewis et al. (2026) for further details.

Objects of interest. There are two objects of interest. The first is the discrete *unconditional RIMP distribution*

$$\{\lambda_g, \pi_g\}_{g=1}^{G}, \tag{4}$$

which is consistently estimated provided Equation (3) is correctly specified. This distribution reveals the different RIMPs and the shares of firm-time observations that are associated with

each. The second object of interest is the posterior *predicted RIMP distribution* for each j,

$$\left\{\hat{\lambda}_j = \sum_{g=1}^{G} \gamma_{j,g}\lambda_g\right\}_{j=1}^{N}, \tag{5}$$

which minimizes the mean squared error of the predicted value of $\Delta \log k_j$. It amounts to the econometrician's optimal prediction of the change in capital for observation $j = (i, t)$ and is a weighted average of λ_g estimates. Examining all estimates of $\hat{\lambda}_j$ allows us to characterize a full distribution of RIMPs across (i, t) observations. This distribution reveals how the response of investment to a monetary policy surprise varies across firms and time periods.

Alternative approaches. Apart from the approach taken by most of the literature, formally presented in Equation (2), there are several alternatives to our procedure for estimating heterogeneity in the RIMP. One option would be to run the specification in Equation (1) separately for each firm. This approach is taken by Aruoba and Drechsel (2024) to study the responses of individual price categories to changes in monetary policy. It is a different way to reduce the dimensionality of the problem, amounting to estimating a λ_i that varies in the cross section, while shutting down time variation entirely. For modeling firms' investment responses to monetary policy, we think modeling time variation is crucial, as in many structural models the RIMP will depend on idiosyncratic state variables that vary over time; our estimates reinforce this intuition. Another alternative would be to follow the approach in Equation (2) but consider "many" possible drivers of heterogeneity x. This approach is taken by Krusell et al. (2023) using machine learning techniques. However, their approach also does not allow for time variation, while our approach does. It further relies on selecting x *ex ante*, even though a large number of x can be selected. Bandeira (2023) proposes "classifier LASSO" as a regularization method to estimate the distribution of heterogeneous responses using local projections in panel data, with an application to the effect of technology shocks on firm-level debt. However, his approach is also tailored to time-invariant responses. Finally, Huang (2021) adapts K-means clustering (see, e.g., Bonhomme and Manresa, 2015) to local projections, with an application to regional house price responses to monetary policy. However, clustering approaches like K-means are likewise not suited to our setting, since they require a long-T panel structure and fixed group membership for any consistency guarantees.

2.3 Ex-post analysis of RIMP drivers

Once we have estimates of $\hat{\lambda}_j$, we can assess the statistical and economic relationship between those estimates and any firm-level variables, x, or fixed effects *ex post*. We focus on traditional candidates from the literature, but also consider a number of novel predictors. In doing so, we can uncover how the RIMP is *individually* correlated with observables, as well as *joint* correlations. Our approach takes the form

$$\hat{\lambda}_{i,t} = \Psi(\tau_{s,t}, X_{i,t-1}; \psi) + \omega_{i,t}, \tag{6}$$

where $\tau_{s,t}$ are industry-time fixed effects for industry s at time t, $X_{i,t-1}$ is an array of time-varying firm-level covariates (lagged by one period), and ψ is a parameter vector. $\Psi(\cdot, \cdot; \cdot)$ is a general nonlinear function. Among our specifications is the linear case, in which Equation (6) is an OLS regression. We also consider nonlinear specifications and machine learning approaches to estimating Equation (6).

When we estimate these *ex post* regressions, we examine in particular the signs of the coefficients for firm characteristics that have been studied in the previous literature. Since we retrieve RIMPs that vary across firms and over time, we can also examine whether nonlinear relationships with firm-level characteristics arise. The R^2 of such regressions is of interest as well. It reveals what fraction of variation in the RIMP can be explained by selected observables, both individually and jointly.

2.4 Practical considerations

Choice of *G*. The total number of groups is an input to the algorithm, but can be selected based on statistical criteria. We set $G = 4$ following minimization of the Bayesian Information Criterion (BIC). We vary G locally in robustness exercises. The correlation between posterior-weighted RIMPs for $G = 3$ and our baseline $G = 4$ is 0.98; for $G = 5$ and our baseline it is 0.99. We also confirm that our ex-post analysis of RIMP drivers is robust to the choice of G.

Standard errors. We compute standard errors in two alternative ways. First, we derive analytical standard errors using the Fisher information. We cluster these standard errors by time, following the recommendation of Almuzara and Sancibrián (2024). Second, we implement a bootstrap procedure. The bootstrap has advantages and disadvantages. The advantage is that we can account for uncertainty in all steps of our analysis, for example coming from the residualization in (3) or from the fact that the left-hand side in (6) is

estimated. The disadvantage is that the bootstrap is computationally expensive. In the text, we report the analytical standard errors and highlight instances in which the bootstrap standard errors differ. These differences are not qualitatively meaningful.

2.5 Estimation for $h > 0$

While most of our analysis focuses on $h = 0$, we also estimate λ_g^h for horizons $h = 1, ..., 7$. This enables us to obtain group-specific impulse response functions (IRFs) up to two years after a given monetary policy surprise. To construct these IRFs, we fix the group membership we estimate for horizon $h = 0$ and then construct group-specific IRFs using weighted panel local projections. When doing so, we restrict the sample to firms that are present at all horizons up to two years.

3 Results

We organize our results into several parts. We begin by presenting the unconditional and predicted RIMP distributions. Then, we examine sources of RIMP variation, including its observable drivers, by linking our estimates to firm-level observables. We finally present findings about the RIMP at longer horizons.

3.1 The RIMP distribution

Table 1 presents our estimate of the unconditional RIMP distribution, the object represented by (4). The top two rows of the table contain the estimates of the RIMP for each group, λ_g, and associated standard errors. We normalize the monetary policy surprises so that these coefficients correspond to a surprise resulting in a 1 pp increase in the Federal Funds Rate. The bottom two rows show the estimates of the probability of belonging to a given group (or shares of (i, t) observations in each group), π_g, with associated standard errors.

The table reveals that investment falls in response to a monetary policy tightening in all four groups. The coefficient for group 4, which corresponds to the smallest responsiveness, implies that the quarterly growth rate of capital falls by 0.5 pp when interest rates rise by 1 pp. Interestingly, while this coefficient is small, it is significantly different from zero. In other words, there is no group that corresponds to zero responsiveness. About a third of firm-time observations belong to this minimally responsive group.

The coefficient in group 1, the group with the largest responsiveness, is about 8 times larger than the estimate for the least responsive group; the quarterly growth rate of capital falls by 4.2 pp when interest rates rise by 1 pp. For a typically sized monetary policy

Table 1: Investment responses to monetary policy tightening – unconditional distribution

	group 1	group 2	group 3	group 4
RIMP (λ_g)	-4.21	-1.66	-0.90	-0.51
Standard Error	(1.81)	(0.35)	(0.09)	(0.04)
Probability of being in group (π_g)	0.05	0.19	0.45	0.32
Standard Error	(0.00)	(0.01)	(0.01)	(0.01)

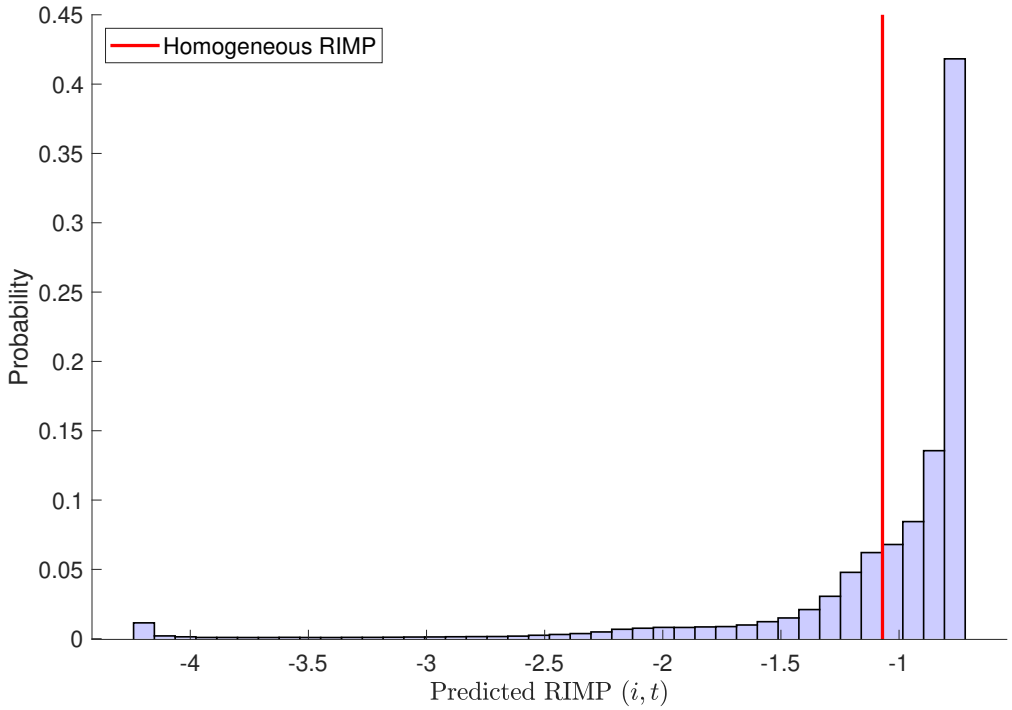
Notes: Estimates of the horizon-0 RIMP and group membership probabilities. Standard errors are in parentheses, based on the Fisher Information clustered by time.

tightening of 25 bp, this corresponds to a fall in the quarterly growth rate of capital of about 1 pp. Only about 5% of firm-time observations are in this "strongly responsive" group. Groups 2 and 3 have coefficients of 1.7 and 0.9, respectively. Almost two thirds of firm-time observations belong to these "moderately responsive" groups.

The estimates of λ_g in Table 1 measure the semi-elasticity of the growth rate of capital to interest rates. Assuming an annual depreciation rate of 10%, as in Ottonello and Winberry (2020), our estimates imply a semi-elasticity of annualized investment to interest rates that ranges between 5% and 42%. Ottonello and Winberry (2020) report 20% as an annual average across all firms. For further discussion, see Zwick and Mahon (2017) and Koby and Wolf (2020). Taken together, the picture emerging from Table 1 is that the investment of most firms in most time periods does not respond strongly contemporaneously ($h = 0$) to monetary policy, while a small number of firms in some periods are much more sensitive to a change in interest rates.

Figure 1 plots the predicted RIMP distribution, the object formally represented by Equation (5). It presents a histogram over firm-time observations, where the magnitude of the RIMP is on the x-axis and the corresponding density is displayed on the y-axis. While Table 1 reports the discrete unconditional distribution of responsiveness, Figure 1 depicts the distribution based on the econometrician's best prediction of the RIMP for each firm-time observation. This prediction is calculated as a weighted average of the four different λ_g estimates following Equation (5). The figure also contains the RIMP estimated from a homogeneous regression (1) in red. This estimate is very similar to the average of the predicted RIMPs across the histogram (they would be identical if the panel were balanced). The histogram reveals that there is a large mass of firm-time observations corresponding to small RIMPs. The stronger the responsiveness (moving from right to left in the graph), the lower is the mass of firm-time observations. The mass starts increasing again around -4, with a minor mode at the maximum responsiveness of around -4.2.

Figure 1: Investment responses to monetary policy tightening – predicted distribution



Notes: The histogram plots the predicted RIMP at the firm-quarter level, as computed by Equation (5). Baseline sample of 157,101 firm-quarter observations.

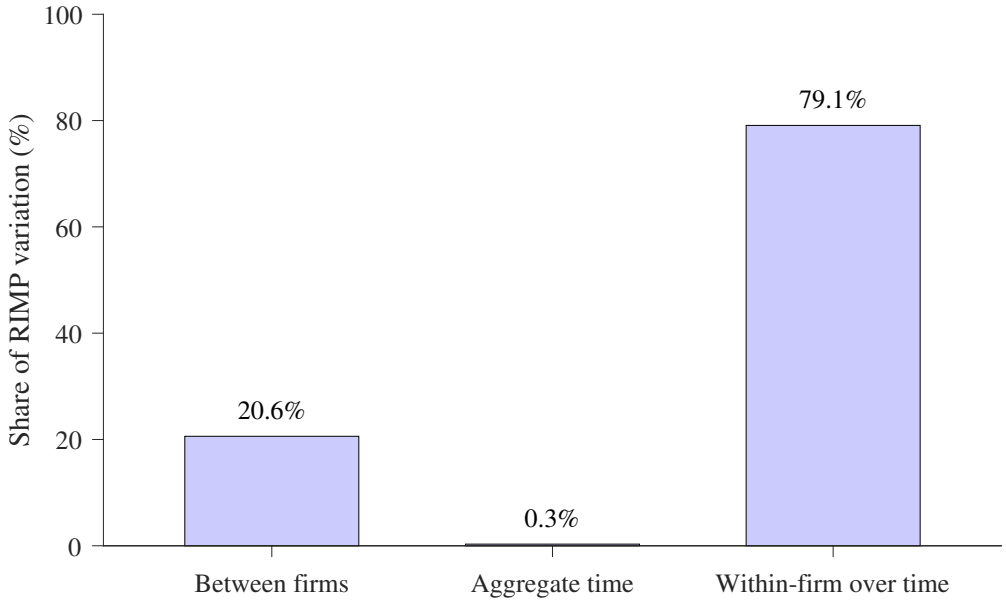
3.2 RIMP variation across firms vs. across time

We begin investigating the underlying sources of RIMP variation. First, we decompose the variance in the posterior predicted RIMP into firm variation, aggregate time variation, and firm-specific time variation. Specifically, we estimate regressions of $\hat{\lambda}_{i,t}$ on quarter fixed effects only, firm fixed effects only, and both sets of fixed effects. We then use a Shapley decomposition that averages each factor's marginal contribution to the R^2 over possible orderings. The result is presented in Figure 2. 20.6% of RIMP variation is attributed to time-invariant differences across firms, only 0.3% to aggregate movements common across firms, and 79.1% to firm-specific variation over time.

This first finding on the sources of RIMP variation implies that the responsiveness to monetary policy is a feature of firms that changes over time, rather than a fixed firm characteristic or a business cycle effect common to all firms. In Appendix A.1, we also show that the RIMP distribution represented in Figure 1 is quite stable through time, and that the average RIMP is not correlated with real US GDP growth. While we do not yet connect RIMP variation to observable firm characteristics, the findings here already indicate that time-varying firm observables might be better candidates for predicting RIMP variation.

Second, we compute a transition matrix, shown in Table 2. We define the "modal RIMP"

Figure 2: Sources of variation in investment responses to monetary policy



Notes: Shapley decomposition of predicted RIMP variation into permanent differences across firms, aggregate variation over time, and within-firm variation over time. The three components sum to 100%.

as the RIMP associated with a firm-time observation's highest posterior probability group. In the table, each row corresponds to a modal RIMP in time t, while each column corresponds to a modal RIMP in time $t+1$. A given cell displays a transition probability between time t and $t+1$. For example, a firm that cuts investment by 0.9 pp after a 1 pp monetary tightening at time t has a 33% probability of cutting investment by 0.5 pp in response to a monetary policy surprise in $t+1$. Note that this is a transition from one impact response to another impact response. The table does *not* show the dynamic responses of investment to a monetary policy change, which would instead consider a single shock and how it plays out over $h > 0$; we consider impulse response estimation separately below.

Above all, Table 2 reveals that strong RIMP status is highly transient. A firm that cuts investment by 1 pp in response to a 25 bp interest rate hike in a given quarter will respond with the same sensitivity in the next period with only 12% probability. A firm that responds weakly, will also respond weakly in the following period with a 59% probability.

Our finding that firms display a sizable RIMP only infrequently suggests that their responsiveness to interest rate changes exhibits a lumpy pattern. While we study investment conditional on the realization of a monetary policy surprise, it is well known that lumpy behavior is also present in firm investment itself (see, e.g., Caballero and Engel, 1999; Baley and Blanco, 2021). This raises the question of how the RIMP and unconditional investment behavior are related. In the next section, we therefore examine the relationship between

Table 2: Investment responses to monetary policy tightening – transition matrix

		modal RIMP in $t+1$				
		λ_1	λ_2	λ_3	λ_4	
modal RIMP in t		-4.21	-1.66	-0.90	-0.51	Total
λ_1	-4.21	0.12	0.23	0.42	0.23	1.00
λ_2	-1.66	0.06	0.29	0.44	0.21	1.00
λ_3	-0.90	0.02	0.10	0.55	0.33	1.00
λ_4	-0.51	0.01	0.05	0.34	0.59	1.00
Total		0.03	0.10	0.45	0.43	1.00

Notes: Each cell reports the transition probability from the time t RIMP in the left column to the time $t+1$ RIMP in the top row. Results are based on the modal RIMP, corresponding to the group with the highest posterior probability for observation j.

the predicted RIMP and lumpiness measures, such as an indicator of investment inaction and the distance to the last investment spike. Beyond that analysis, a potential concern may be that our estimation spuriously assigns observations to different RIMP groups only due to variation in firms' unconditional investment behavior. In an alternative specification, we therefore allow for group-specific intercepts α_g and repeat our estimation. We find that the estimates of $\hat{\lambda}_{i,t}$ we obtain from that specification have 0.99 correlation with those in our baseline specification. In other words, allowing for a group-specific intercept does not materially affect our findings. Results for the specification with group-specific intercepts are presented in Appendix A.4.

3.3 Observable drivers of the RIMP distribution

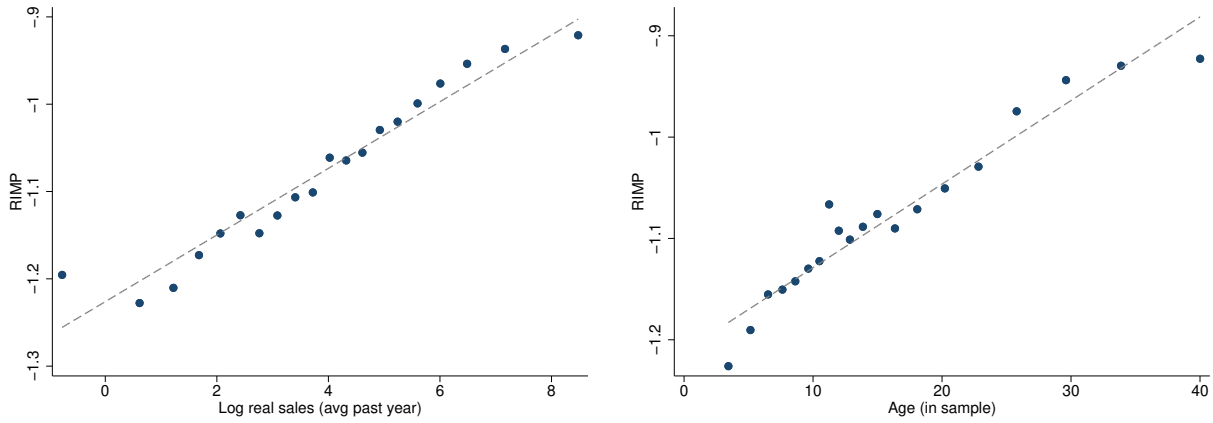
We next examine the role of observable firm characteristics in explaining firm heterogeneity in the investment channel of monetary policy. We do so by relating our estimates to firm-level characteristics *ex post*, estimating Equation (6). That is, we regress the predicted RIMPs on prominent candidates from the literature, such as size, age, distance to default, or debt maturity. We also consider novel drivers from recent influential work in corporate finance.

3.3.1 Traditional firm characteristics and the RIMP

One of the classic explanations for differences in the responsiveness to monetary policy is firm size, following for example Gertler and Gilchrist (1994). Figure 3a plots a binscatter where size, measured as the log of real sales averaged over the past year, is on the x-axis

and our predicted RIMP estimate is on the y-axis.[2] Smaller firms have a larger RIMP in absolute value, that is, a larger decline in investment after an interest rate increase. This confirms the insights from the literature. It is consistent with firm size capturing financial frictions, for instance. The patterns in Figure 3a are virtually unchanged when we consider the log of real total assets as a measure of size, or when we use the simple lag rather than the 4-quarter average of the size measure. In Figure 3b, we repeat the same analysis with firm age and find that younger firms are more sensitive to monetary policy, confirming results by Cloyne et al. (2023).

Figure 3: RIMP heterogeneity by firm size and age



(a) RIMP and size

(b) RIMP and age

Notes: The figure presents binscatter plots based on 20 quantiles. The predicted RIMP is plotted against size (left, the log of average real sales over the past four quarters) and age (right, defined as years in sample). Dashed lines indicate fitted linear regressions. Regressing the RIMP on size, we obtain a coefficient of 0.038, significant at the 1% level. Regressing the RIMP on age, we obtain a coefficient of 0.008, significant at the 1% level.

Figure 4 exhibits additional characteristics from the literature. The first row of the figure presents binscatter plots for leverage and distance to default, following Ottonello and Winberry (2020). OW find that more leveraged firms are less responsive to monetary policy. When computing the linear correlation between RIMP and leverage, we confirm their result. However, Figure 4a reveals nonlinearities in this relationship. We find that firms with leverage that is either much lower or much higher than their average are the most responsive to monetary policy. With regard to distance to default, we find an ambiguous relationship, shown in Figure 4b, whereas OW find that less risky firms (i.e. those with higher distance to default) are more responsive to changes in monetary policy. We discuss the possible reasons

[2]In this section, we visualize the relationship between the RIMP and predictors via quantile-based binscatter plots and fitted curves. In Appendix Figure A.3 we use an alternative approach, binsreg, discussed by Cattaneo et al. (2024).

behind this disconnect below.

In Figure 4c, we show that firms with higher cash-to-assets ratios are more responsive to monetary policy. Jeenas (2023) finds the opposite, but only at longer horizons: upon impact, he finds an insignificant coefficient on the interaction term (i.e. no correlation). Consistent with the channels proposed by Jeenas (2023), we find that large debt issuance is associated with a larger (absolute) RIMP. Finally, we show in Figure 4d that firms with shorter debt maturity are more sensitive to monetary policy, confirming results by Jungherr et al. (2024).[3]

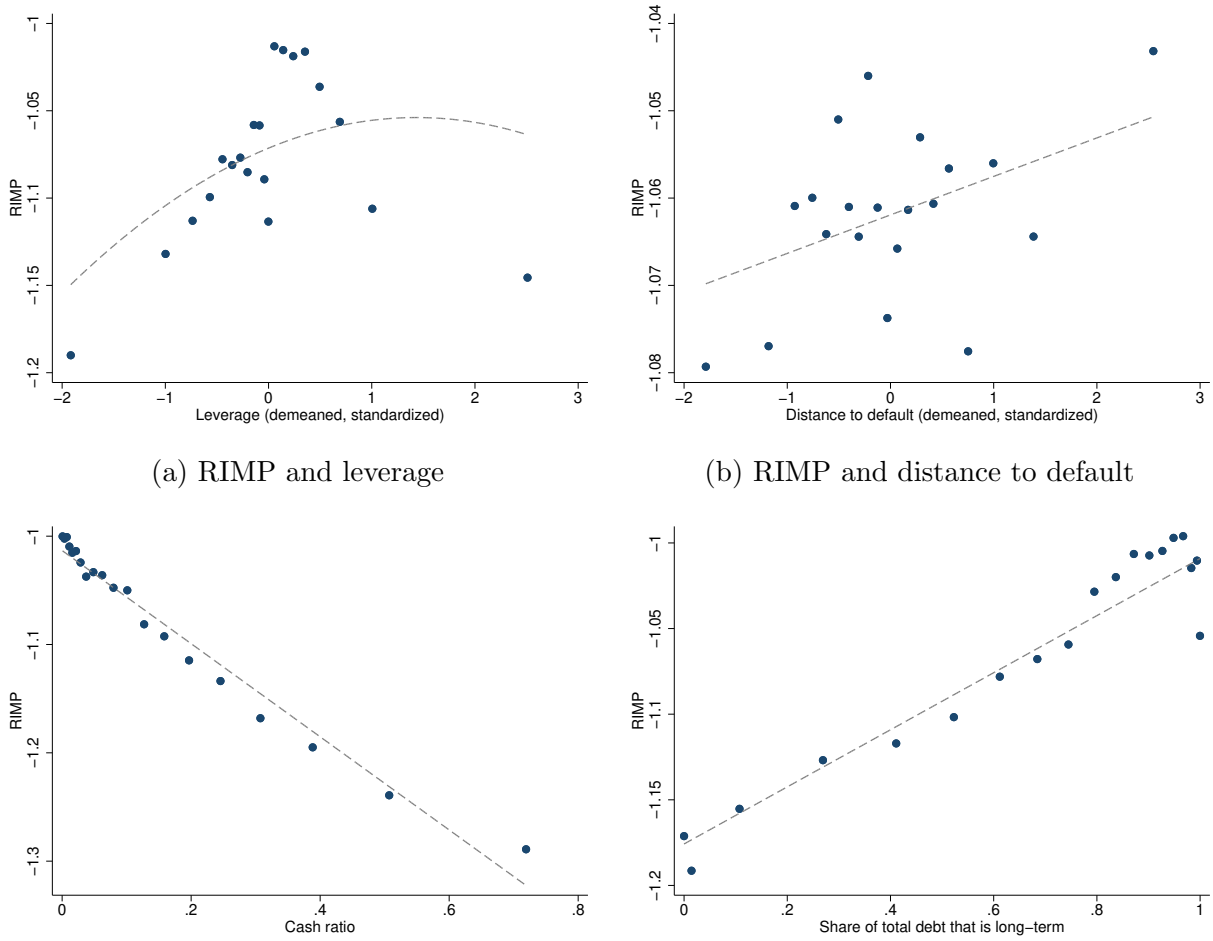
There are several reasons why the interaction term approach followed by the previous literature may imply a different relationship from the correlation between the RIMP and a given firm characteristic.[4] First, the "true" heterogeneity in the RIMP could be multi-dimensional, and thus be driven by factors other than those captured by a selected observable characteristic such as distance to default. A univariate interaction ignores these other factors, potentially leading to omitted variable bias. We provide an example in Appendix B. In the interaction approach that is standard in the literature, such concerns could be addressed by saturating the regression with many interactions, although this would come at the cost of statistical power. Second, the "true" heterogeneity in the RIMP could be highly nonlinear. Even if the right variable is included through an interaction term, a linear interaction would not correctly capture RIMP heterogeneity. David et al. (2026) study this form of misspecification of local projections in detail. We further examine nonlinearities in the RIMP below. A third reason is that a substantial portion of heterogeneity in the RIMP could be driven by *unobserved* factors; these will always be omitted from strategies exploiting interactions. Indeed, we find that observable drivers explain only a small fraction of overall heterogeneity in the RIMP. One way to see this is to consider the economic magnitudes in Figures 3 and 4. We have documented that the RIMP ranges from -0.5 to -4.2. However, in Figure 3a, the RIMP variation that is associated with firm size, for instance, is between -0.85 and -1.2. For other drivers, such as leverage, the explained dispersion is even smaller. The predicted difference in RIMP between having only long-term debt and no long-term debt is about 0.15 pp, less than half of the inter-quartile range of RIMPs that we estimate.

Even when considered jointly, observable characteristics continue to explain a small fraction of the overall RIMP heterogeneity, as we show in Table 3. In column (1), we regress the RIMP on several firm-level characteristics, while in column (2) we also control

[3]We also examine debt maturity demeaned at the firm level, as in Deng and Fang (2022). In this case we find a non-linear, inverse V-shaped relationship, similar to the patterns we detect for leverage.

[4]We have confirmed that the discrepancies are not due to differences in the sample or in the specification used. In our sample, when we repeat the same regression estimation as in Ottonello and Winberry (2020) (Table 3 in their paper) we confirm their results. Likewise, we confirm their results using linear interactions of their x variables with the orthogonalized monetary policy surprise we use.

Figure 4: RIMP and different firm characteristics



(a) RIMP and leverage

(b) RIMP and distance to default

(c) RIMP and cash ratio

(d) RIMP and debt maturity

Notes: The figure presents binscatter plots based on 20 quantiles. The predicted RIMP is plotted against the following variables: in the top left panel, leverage (the ratio of lagged total debt ($dlcq+dlttq$) and lagged total assets atq, where abbreviations are Compustat variable codes) winsorized and demeaned at the firm level as in Ottonello and Winberry (2020); in the top right, distance to default, winsorized and demeaned, computed by Ottonello and Winberry (2020); in the bottom left, cash ratio (the ratio of lagged cash $cheq$ and equivalent over lagged total assets atq); and in the bottom right, debt maturity constructed as in Deng and Fang (2022) (the lagged ratio of debt maturing in more than 1 year ($dlttq$) over total debt ($dlcq+dlttq$)). Dashed lines indicate fitted regression lines. For leverage, a quadratic specification yields coefficients of 0.024 and -0.009 on the linear and squared terms, both significant at the 1% level. Regressing the RIMP on distance to default yields a coefficient of 0.004, significant at the 1% level; on cash ratio, a coefficient of -0.431, significant at the 1% level; and on debt maturity, a coefficient of 0.167, significant at the 1% level.

for industry-time fixed effects, and in column (3) we include firm fixed effects. When tested jointly, the statistical association between firm-level predictors and the RIMP remains. However, the R^2, a measure of the share of the overall RIMP heterogeneity explained by predictors, is only around 4% in the absence of any fixed effects. Adding industry and

Table 3: RIMP and observable characteristics - multivariate analysis

	(1)	(2)	(3)
Leverage (demeaned, standardized)	-0.001	-0.002	0.002
	(0.003)	(0.003)	(0.003)
Distance to default (demeaned, standardized)	0.004**	0.005**	0.000
	(0.002)	(0.002)	(0.002)
Share of total debt that is long-term	0.096***	0.098***	0.035***
	(0.006)	(0.006)	(0.008)
Cash ratio	-0.330***	-0.342***	-0.142***
	(0.014)	(0.014)	(0.027)
Firm age	0.004***	0.004***	0.001**
	(0.000)	(0.000)	(0.001)
Log of real sales (avg past year)	0.022***	0.023***	0.017***
	(0.001)	(0.001)	(0.005)
Observations	110395	110379	110307
R-squared	0.041	0.053	0.206
Adjusted R-squared	0.041	0.049	0.172
Time x Industry FE		✓	
Firm FE			✓

Notes: In each column the dependent variable is the predicted RIMP at horizon 0, estimated with baseline time-varying group membership. Regressors are lagged by one quarter. Robust standard errors are reported in parentheses. *,**, and *** denote significance at the 10, 5, and 1% levels respectively.

time heterogeneity marginally increases the R^2 to close to 5%. However, the RIMP does in fact vary by industry, with firms in service-oriented industries (i.e. SIC 2-digit codes 70-89) responding the most and agricultural firms the least (see Appendix Figure A.4). Nevertheless, the extent of this variation is economically and statistically small. Finally, adding firm fixed effects increases the R^2 to almost 21%, implying that the majority of RIMP variation remains unexplained after absorbing variation in the firm-specific average RIMP.

Economically, Table 3 leads to similar conclusions as our previous variable-by-variable analysis, for most of the firm observables. Size, age, liquidity and the long-term debt share have the same signs as implied by the bin scatter plots in Figures 3 and 4. Interestingly, the coefficient on leverage becomes statistically indistinguishable from zero after controlling for the other variables. Distance to default is also associated with large standard errors. In

fact, it becomes statistically insignificant when we consider bootstrap-based standard errors. This connects to our discussion of leverage and distance to default above in our discussion of Figures 4a and 4b.

The right-hand-side variables included in Table 3 are those that the recent literature on monetary policy and firm heterogeneity focuses on. That literature is the one that heavily uses panel local projections with univariate interactions, motivating our new approach. In addition to these variables, we construct several firm observables that are motivated by the broader firm dynamics and investment literature. That literature has put a stronger focus on real investment frictions, productivity heterogeneity and misallocation. The measures we examine include the market-to-book ratio, an indicator of investment inaction, an indicator of investment spikes, the distance to the last investment spike, a proxy for firms' total factor productivity (TFP) and a proxy for the marginal revenue product of capital (MRPK).[5] As shown in the Appendix Table A.1, these measures do not substantially increase the R^2.

Furthermore, we run a LASSO regression that includes a long list of variables and allows for nonlinearities and interactions across all variables, amounting to more than 200 regressors. The LASSO fails to increase the R^2 above 9% for the specification that includes time-industry fixed effects corresponding to column (2). The strongest predictors chosen by the LASSO include some of the characteristics previously discussed, such as age, size, and the share of debt that is long-term, as we show in Appendix Figure A.5.[6]

3.3.2 Implications of unobservable RIMP heterogeneity

There are at least three ways to understand the limited explanatory power of typical firm-level observables in the context of empirical research on monetary policy and firm heterogeneity. First, there is no clear benchmark for how high the explanatory power of firm characteristics for the RIMP should be. Heterogeneity in household MPCs is likewise largely unexplained by observable characteristics (Lewis et al., 2026). There is also large unobserved heterogeneity in unconditional investment behavior. If we replace the RIMP with investment itself ($\Delta \log k_{i,t}$) in the regressions underlying Table 3, the corresponding R^2 in columns (1), (2) and (3) would be 0.027, 0.044, and 0.117.

Second, unobserved heterogeneity in the RIMP may be less relevant if a researcher is interested in the aggregate RIMP. The portion of micro-level RIMP variation that is explainable with observable firm-level characteristics might drive the response of aggregate

[5] In a recent study, González et al. (2023) include a measure of the MRPK as an interaction term with a monetary policy surprise in a panel local projection, specified as in Equation (2). TFP and MRPK are of course subject to measurement challenges, as shown for example by Bils et al. (2021).

[6] In the Appendix, we also verify that the conclusions from Table 3 do not change when we vary the number of groups G, see Table A.3.

Table 4: Alternative RIMP aggregations

	Implied aggregate RIMP
Using headline RIMP estimates	-0.909
Using RIMP estimates predicted by leverage only	-1.062
Using RIMP estimates predicted by size only	-0.895
Using RIMP estimates jointly predicted by observables	-0.896

Notes: Aggregate RIMPs are computed from $\hat{\lambda}_{i,t}$ estimates weighted by each firm's share of the total capital stock. The first line uses our headline estimates, $\hat{\lambda}_{i,t}$, while subsequent lines aggregate the predicted values of these RIMPs from a regression on the indicated firm characteristic(s).

investment to monetary policy, in particular when the RIMP of large firms is well captured. To illustrate this idea, we calculate the implied aggregate RIMP from different micro-level RIMP estimates. For any set of estimated $\hat{\lambda}_{i,t}$, we can compute an implied aggregate $\hat{\lambda}_{agg}$ by summing the $\hat{\lambda}_{i,t}$, weighted by each firm's share in the total capital stock.[7] We can construct this aggregation separately for the headline RIMPs that we estimated, or for the predicted values obtained from a projection on observables, such as size and leverage.

Table 4 presents aggregations of the response to a 1 pp monetary policy surprise using different micro-level RIMP estimates. Using headline RIMP estimates gives an aggregate RIMP of -0.909, as shown in the first line of the table. This is 15 bp larger (i.e., closer to zero) than the homogeneous RIMP estimate in Figure 1 because small firms are more responsive to monetary policy surprises. The following lines of the table aggregate only that RIMP variation that can be explained by leverage, size, or both. All implied aggregated RIMPs are fairly close to the same estimate, ranging from -1.062 to -0.895. This shows that the unexplained micro-level heterogeneity is not required to capture the aggregate macro-level RIMP. Size, in particular, is sufficient to approximate the aggregate response.

Finally, it should be noted that the unobserved heterogeneity we detect should motivate new research to further uncover the investment channel of monetary policy, in particular by considering new firm observables. We provide some new results in this dimension in the following subsection.

3.3.3 Novel firm characteristics affecting the RIMP

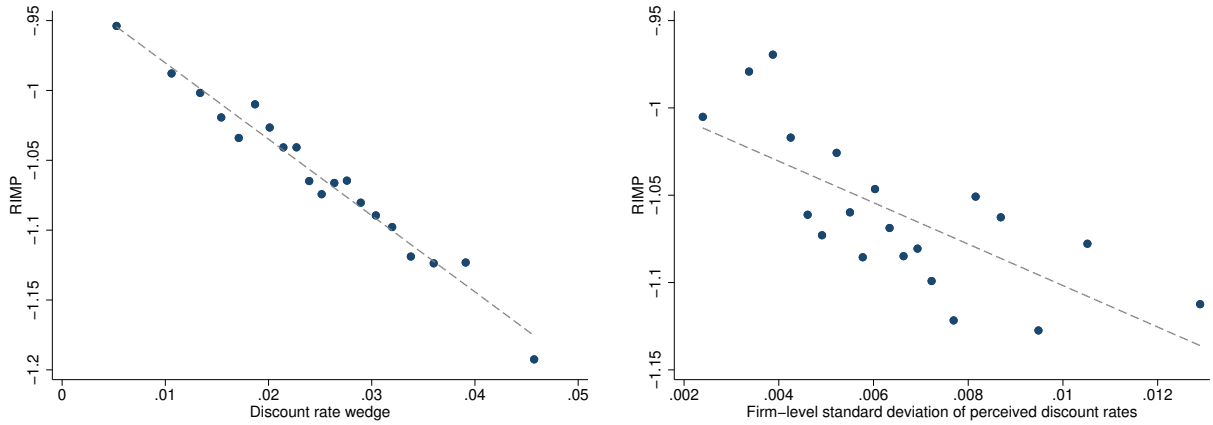
We now turn to recent influential studies on corporate decision making in search of novel drivers that may explain a greater share of variation in the RIMP. These studies go beyond the information about firms' financial statements and balance sheets that can be found in

[7] Formally, we use the fact that $\hat{\lambda}_{agg,t} \equiv \frac{\partial \log \sum_{i,t} k_{i,t}}{\partial \epsilon_t^m} = \sum_{i,t} \frac{k_{i,t-1}}{\sum_{i,t} k_{i,t-1}} \hat{\lambda}_{i,t}$, before also averaging across time to obtain $\hat{\lambda}_{agg}$.

Compustat by measuring firms' characteristics and decisions through processing corporate earnings calls and conducting comprehensive surveys of firm leadership personnel.

Gormsen and Huber (2025) construct measures of firms' perceived cost of capital and perceived discount rates from corporate earnings calls. They document a significant wedge between these two objects, which is strongly correlated with firm-level investment. This finding raises the question of whether firms' assessments of funding costs and discount rates also determine their responsiveness to changes in the general level of interest rates, as induced by monetary policy. Using the publicly available version of the data from Gormsen and Huber (2025), we find that the RIMP is strongly negatively correlated with these behavioral wedges.[8] We document this fact in Figure 5a, which shows a more strongly negative RIMP for higher discount rate wedges. This correlation remains strong and statistically significant when we control for the regressors reported in the first column of Table 3.

Figure 5: RIMP heterogeneity by discount rate wedges and stickiness



(a) RIMP and discount rate wedges

(b) RIMP and (inverse) discount rate stickiness

Notes: The figure presents binscatter plots based on 20 quantiles. The predicted RIMP is plotted against the discount rate wedge, the difference between perceived discount rates and perceived cost of capital (left panel) and the standard deviation of perceived discount rates over time (right panel). The data come from Gormsen and Huber (2025), are available at quarterly frequency since 2002, and yield a sample of 34,952 observations. Dashed lines indicate fitted linear regressions. Regressing the RIMP on discount rate wedge yields a coefficient of -5.5, significant at the 1% level and on the standard deviations yields a coefficient of -11.9, significant at the 1% level.

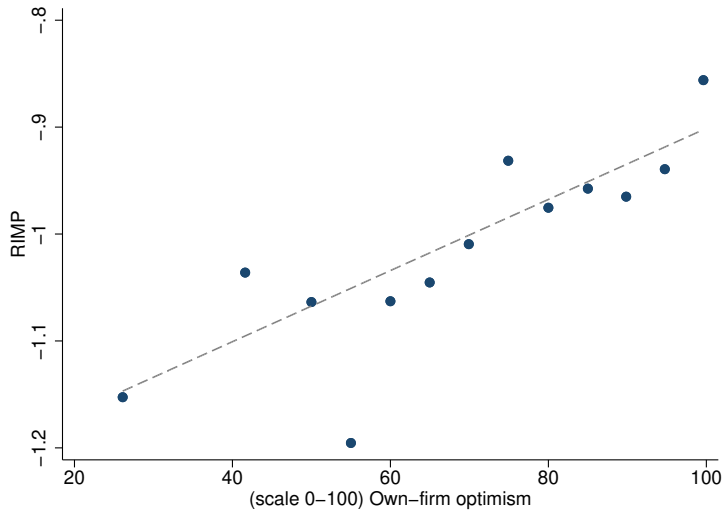
In additional research on corporate discount rates, Fukui et al. (2024) document that perceived discount rates are insensitive to expected inflation, interpreting this behavior as

[8]While Gormsen and Huber (2025) measure perceived discount rates and cost of capital directly at the firm level, the publicly available data we use consists of predicted values that the authors compute based on observable characteristics for a wider set of firms. Therefore, we see these measures as a conservative estimate of the true unobserved variation in perceptions. We thank Kilian Huber and Niels Gormsen for making their estimates publicly available.

"stickiness" of discount rates. In their structural model, stickier discount rates imply weaker investment elasticity to monetary policy tightening. In the data, we proxy this concept by computing the firm-by-firm standard deviation of perceived discount rates from Gormsen and Huber (2025). As we show in Figure 5b, firms with more volatile discount rates (i.e. less sticky discount rates) have a more negative RIMP, consistent with the channels in Fukui et al. (2024). As for the analysis of discount rate wedges, the relationship remains strong and statistically significant when controlling for other firm-level predictors.

Firms can differ in terms of their overall economic outlook in ways that are not reflected in income statements or balance sheet information available through Compustat. These expectations could plausibly be correlated with the investment sensitivity to monetary policy. To investigate this idea, we use data from the Duke CFO Outlook Survey (Graham and Harvey, 2020) merged with Compustat at the firm level.[9] The RIMP may also operate through firms' adjustment of long-horizon investment plans to monetary policy (see, e.g., Selgrad and Siani, 2025, for a recent study on the topic).

Figure 6: RIMP heterogeneity by CFO optimism



Notes: The figure presents a binscatter plot based on 13 quantiles. The predicted RIMP is plotted against a self-reported measure of CFO optimism towards their own firm, on a scale between 0 and 100. The data come from a vintage of the CFO Business Outlook Survey and are merged with Compustat using the anonymized PERMNO identifier. The data are available since 2001; the merged dataset consists of 692 observations. The dashed line indicates a fitted linear regression. Regressing the RIMP on optimism yields a coefficient of 0.003, significant at the 1% level.

We find that the RIMP is significantly correlated with self-reported expectations. Specifically, the survey asks CFOs to rate their optimism about the financial prospects

[9]We thank John Graham for kindly providing the data.

of their own firm on a 0–100 scale. As shown in Figure 6, firms with less optimistic outlooks about their own prospects are more responsive to monetary policy. The relationship remains statistically significant when we control for other firm-level predictors.[10]

3.3.4 The multidimensionality of the RIMP

A key advantage of our approach is that we can easily visualize how the RIMP changes with multiple observable drivers. Figure 7a shows that small *and* young firms are the most investment-elastic to monetary policy. Several papers have highlighted the joint importance of size and age for firm dynamics (see, e.g., Haltiwanger et al., 2013). These firms are often regarded in the literature as the most financially constrained. This interpretation, however, contrasts with some of our other findings. For example, in Figure 7b, we focus on two key state variables in heterogeneous firm models: size (defined as lagged log capital) and the net leverage ratio (leverage ratio minus cash ratio). We show that small and less leveraged firms are the most investment-elastic to monetary policy surprises. Visually, firm size seems to be much more important than the net leverage ratio.[11]

Finally, it is also possible that firm characteristics are more informative for certain segments of the RIMP distribution. We assess this through a fractional multinomial logistic regression of $\gamma_{j,g}$, the posterior probability of observation j belonging to group g, on firm characteristics. The largest effects are typically observed for the group with the RIMP closest to 0. For instance, a 1 pp higher cash ratio lowers the probability of being investment-inelastic by 0.18 pp, while it increases the probability of having any of the other, more negative, RIMPs, albeit by a smaller amount. The effects of size and debt maturity seem to be more uniformly distributed across groups, as we show in Appendix Table A.2. Importantly, these effects remain statistically significant across most groups, suggesting that the correlations we find are not local. Finally, the McFadden Pseudo-R^2 suggests that the tail (i.e. largest magnitude) RIMP is the easiest to predict with the set of firm observables we consider.

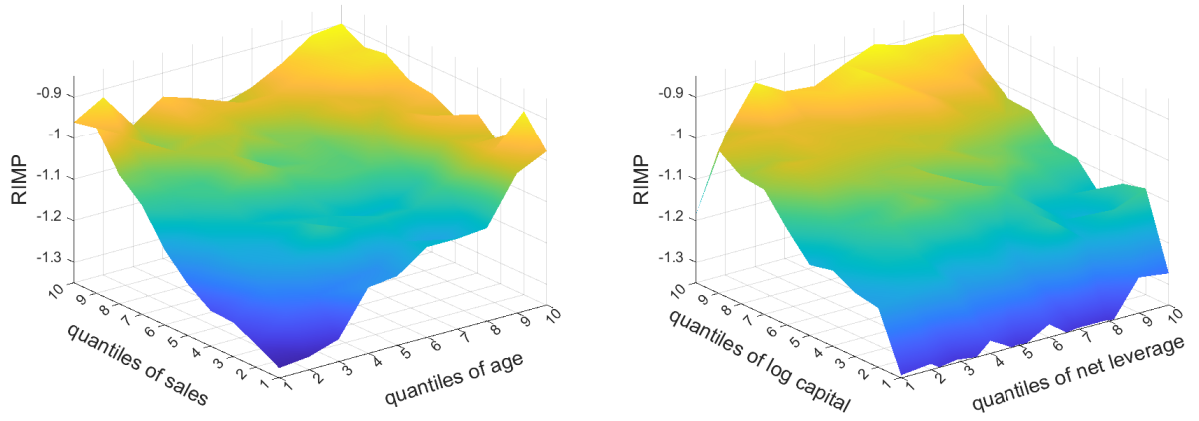
3.4 The RIMP at longer horizons

So far, our analysis has focused on the impact response of firm investment to a change in monetary policy. This focus ensures comparability with the literature, where many papers only report impact responses. Since monetary policy is often thought to operate with long

[10]The survey also asks firms' leadership to rate their optimism towards the U.S. economy as a whole. We find a similar but weaker correlation for this variable, which is statistically insignificant.

[11]Nevertheless, the association between the RIMP and the net leverage ratio remains statistically significant even when we control for lagged log capital.

Figure 7: RIMP multidimensionality



(a) RIMP by size and age

(b) RIMP by capital and net leverage ratio

Notes: In the left panel, the predicted RIMP is plotted against deciles of sales (log sales over the past year) and age. In the right panel, the predicted RIMP is plotted against deciles of lagged log capital and the lagged net leverage ratio (as in Ottonello and Winberry, 2020, we sum lagged current total liabilities and lagged long-term debt, subtract lagged total current assets, and divide by lagged total assets).

and variable lags, it is natural to also study the responsiveness of investment to changes in monetary policy at longer horizons.

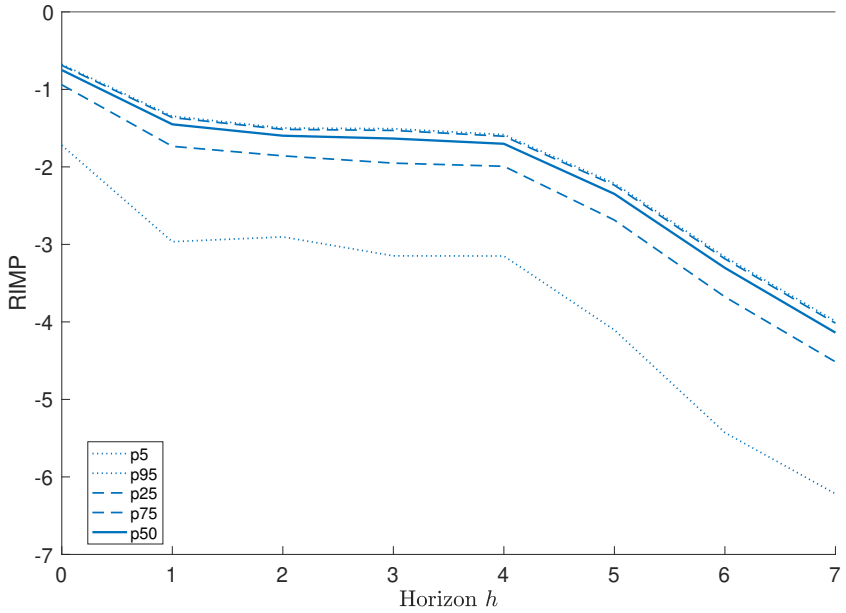
Figure 8 examines heterogeneity in the RIMP for horizons up to two years. Specifically, we fix the group membership we estimate for horizon $h = 0$ and then simply construct group-specific IRFs using weighted panel local projections. To visualize the results, at each horizon $h = 0, 1, ..., 7$ we compute percentiles of the corresponding predicted RIMP distribution. In other words, each vertical slice of Figure 8 presents percentiles of a distribution such as the one shown in Figure 1, but for different horizons.[12]

The IRF distribution shows that there is sizable and growing heterogeneity at longer horizons. The skewness of the RIMP distribution, with few observations corresponding to the largest RIMP in absolute value, appears to be quite persistent. Two years after a monetary policy surprise, the vast majority of firms display an economically meaningful investment response. This is different from the $h = 0$ picture, where the smallest responses are statistically significant, but economically not very large.

As we did for the impact responses, we correlate the predicted RIMP at different horizons with different firm characteristics. The results are similar to what we found before, with small, young, and liquid firms being the most elastic. When we project the RIMP at different horizons on firm predictors jointly, our conclusions are also little affected. Table 5 shows

[12]The horizon-0 responses differ from Figure 1 because the sample used to construct Figure 8 is slightly different, since we only include firms that are present in the sample up to horizon $h = 7$.

Figure 8: Predicted RIMP: impulse responses



Notes: We compute the predicted RIMP ($\hat{\lambda}_{i,t,h} = \sum_{g=1}^{G} \gamma_{i,t,g}\lambda_{g,h}$) for each horizon, and then plot the quantiles (5th, 25th, 50th, 75th, and 95th percentiles) of $\hat{\lambda}_{i,t,h}$ across firms at each horizon. Sample size is restricted to firms with a non-missing dependent variable for all horizons.

that correlations are quantitatively similar and slightly stronger at longer horizons, whereas the adjusted R^2 does not change meaningfully.

4 Implications for structural models

This section discusses the implications of our results for structural models of firm heterogeneity and monetary policy. We begin with some broader observations and then confront a specific structural model from the literature with our empirical findings.

4.1 General considerations for the RIMP and structural models

In the literature on household heterogeneity, it is now standard to show the model-generated distribution of MPCs and how the MPC varies with household characteristics (see, e.g., Auclert, 2019; Kaplan and Violante, 2022). We hope that our results enable a similar practice using the RIMP for models of monetary policy and firm heterogeneity. Many prevalent heterogeneous firm models are quite different from each other, as they are often tailored towards the particular firm observable that a given paper chooses to focus on, such as distance to default, liquidity, or debt maturity. Nevertheless, in each of these models,

Table 5: RIMP and observable characteristics - multivariate analysis at longer horizons

	Horizon								
	0			3			7		
Leverage (demeaned, standardized)	0.002	0.002	0.003	0.003	0.002	0.004	0.005	0.004	0.006
	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.004)	(0.006)	(0.006)	(0.006)
Distance to default (demeaned, standardized)	0.003**	0.003**	0.000	0.004**	0.005**	0.001	0.006*	0.006	0.000
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)	(0.004)	(0.003)
Share of total debt that is long-term	0.066***	0.067***	0.020***	0.104***	0.106***	0.031***	0.150***	0.154***	0.052***
	(0.004)	(0.004)	(0.006)	(0.007)	(0.007)	(0.009)	(0.011)	(0.011)	(0.015)
Cash ratio	-0.237***	-0.244***	-0.106***	-0.379***	-0.390***	-0.168***	-0.536***	-0.558***	-0.270***
	(0.009)	(0.010)	(0.018)	(0.015)	(0.015)	(0.029)	(0.024)	(0.025)	(0.048)
Firm age	0.003***	0.003***	0.001***	0.005***	0.004***	0.002***	0.006***	0.006***	0.003***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.001)
Log of real sales (avg past year)	0.015***	0.016***	0.011***	0.025***	0.026***	0.017***	0.037***	0.038***	0.030***
	(0.001)	(0.001)	(0.004)	(0.001)	(0.001)	(0.006)	(0.002)	(0.002)	(0.009)
Observations	102508	102492	102508	102508	102492	102508	102508	102492	102508
R-squared	0.046	0.059	0.222	0.047	0.060	0.223	0.037	0.050	0.194
Adjusted R-squared	0.046	0.054	0.187	0.047	0.055	0.187	0.037	0.045	0.157
Time × Industry FE		✓			✓			✓	
Firm FE			✓			✓			✓

Notes: In each column we regress the horizon-specific predicted RIMP on the listed set of predictors. Regressors are lagged by one quarter. Robust standard errors are reported in parentheses. *,**, and *** denote significance at the 10, 5, and 1% levels respectively.

multidimensional correlations with the RIMP can be computed and disciplined with the empirical counterparts we provide. Importantly, every model implies a full RIMP distribution irrespective of the specific firm features that are central to each model. Our procedure delivers a target for that full RIMP distribution. Moreover, we show below that this target can be critical in selecting an appropriate model.

There are two conceptual differences between the MPC and the RIMP that we wish to stress. First, the RIMP is a general equilibrium object. As such, it is helpful not only for disciplining the firm problem, but also for how that problem interacts with the rest of the economy. In contrast, the MPC is often used as an indirect partial equilibrium target that disciplines monetary policy transmission in HANK models. To be clear, heterogeneous firm models could also be disciplined with partial equilibrium elasticities; see Winberry et al. (2025) and Melcangi (2025) for discussions on sensitivities to cash flow shocks, and Zwick and Mahon (2017), Koby and Wolf (2020), and Wroblewski (2025) for partial equilibrium investment-interest rate elasticities. However, empirical evidence on these firm-level partial equilibrium sensitivities is in scarce supply, making estimates of the RIMP useful.

Second, the RIMP is specific to monetary policy, while the MPC is out of any transitory income change. To the extent that our empirical results on the RIMP can be used to distinguish primitive model elements, such as adjustment costs or financial frictions, they

are relevant to an even broader set of macroeconomic models in which firm investment is crucial, but which do not focus on monetary policy. This includes neoclassical models of optimal firm investment and Q-theory, such as Jorgenson (1963), Hayashi (1982), Abel (1983), and Abel and Eberly (1994), models of fixed adjustment costs and (S, s)-behavior, such as Caballero and Engel (1999) and Baley and Blanco (2021), or models of financial frictions, such as Kiyotaki and Moore (1997) and Khan and Thomas (2013).

Our results offer a new set of empirical targets to discipline structural models featuring firm heterogeneity and monetary policy. We hope they will spur research aimed at understanding how monetary policy transmits through firms. In the following section, we demonstrate how one of our results in particular, the shape of the full RIMP distribution, can discipline models above and beyond what the conventional linear-interaction approach achieves.

4.2 Disciplining a structural model based on our results

We consider the model developed in the recent survey paper on heterogeneity in macroeconomics by Winberry et al. (2025). This heterogeneous firm model is a simplified version of the firm dynamics structure of Khan and Thomas (2013), but with a New Keynesian layer added, as in Ottonello and Winberry (2020). In brief, firms are *ex post* heterogeneous due to idiosyncratic risk to their productivity and are subject to exogenous exit shocks. They produce using capital and labor, the former subject to quadratic adjustment costs and the latter chosen frictionlessly on the spot market. Firms cannot issue equity, whereas they can issue debt, but borrowing is subject to a collateral constraint. We outline the model in Appendix C. To obtain a RIMP distribution, we consider a transitory and unexpected monetary shock, using Sequence Space Jacobian methods.[13] Then, for each firm, we compute the RIMP from simulated data with and without the shock, and normalize it as in our empirical analysis: the change in the capital growth rate induced by a 1 pp hike in the nominal interest rate.

We demonstrate the usefulness of our empirical results for developers of structural models using a specific calibration experiment. The experiment compares two separate calibrations of the model, shown in Table 6. In both models, we fix all parameters at values chosen in the literature, except for those that govern capital adjustment costs (ϕ) and the collateral constraint (χ). In particular, the parameters that determine firms' productivity process, production function, and the initial capital stock are set to the same values as in Winberry

[13] Winberry et al. (2025) (Appendix B) define interest rate sensitivities to a one-time change in the real rate. We thank the authors for making that code available to us. We extend that framework to analyze a monetary shock in general equilibrium.

et al. (2025). The parameters related to the New Keynesian block of the model are set to the same values as in Ottonello and Winberry (2020).

We then choose χ and ϕ such that each of the two models matches the correlation between the RIMP and firm size ($\log k_{i,t-1}$). In Version A, we set the parameter related to the capital adjustment costs to $\phi = 0.20$ and the parameter in the collateral constraint to $\chi = 0.20$. In Version B, we set $\phi = 0.05$ and $\chi = 0.70$. The table notes provide more details about how these parameters enter the model.

Table 6: Alternative versions of the structural model

Parameter	Version A	Version B
Capital adjustment cost parameter (ϕ)	0.20	0.05
Collateral constraint parameter (χ)	0.20	0.70

Notes: We calibrate a general equilibrium version of the model of firm heterogeneity and monetary policy from Winberry et al. (2025). In their paper, convex capital adjustment costs are given by the term $\frac{\phi}{2}\left((k_{j,t} - k_{j,t-1})/k_{j,t-1}\right)^2 k_{j,t-1}$. The collateral constraint is written as $b_{j,t} \leq \chi k_{j,t}$, where $b_{j,t}$ is a one-period non-contingent bond. The other model parameters are calibrated following Winberry et al. (2025) and Ottonello and Winberry (2020). See Appendix Table C.1 for details of all model parameters.

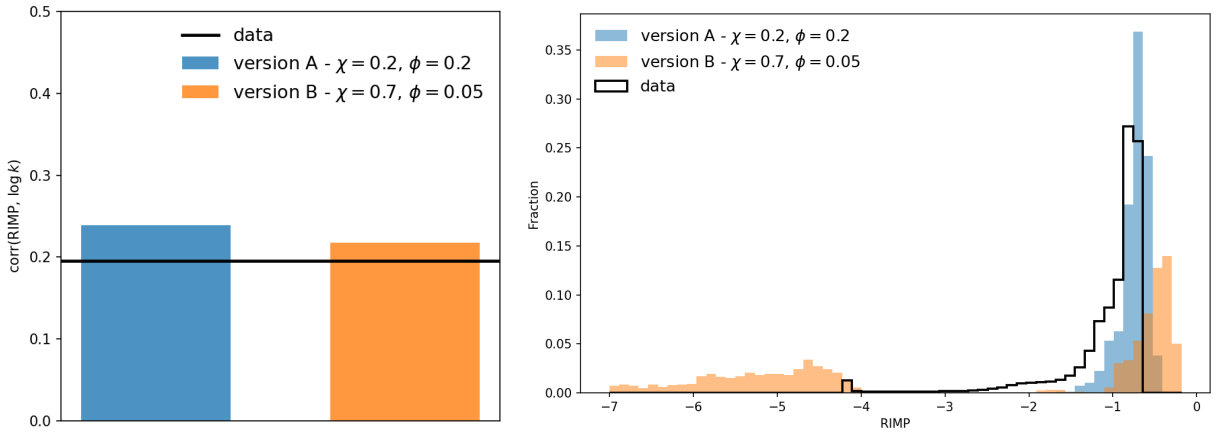
Panel (a) of Figure 9 confirms that both calibrations match the empirical relationship between firm size and the RIMP. The correlation between $\hat{\lambda}_{i,t}$ and $\log k_{i,t-1}$ is 0.2 in the data, and reasonably close to 0.2 in either of the model calibrations. This proximity of the RIMP-size correlation implies that an approach that disciplines structural models using the univariate empirical connection between firm size and the RIMP would accept both calibrations as consistent with the empirical evidence.[14] While we focus on firm size here, the conclusions from this discussion also apply using correlations with other firm characteristics.

Panel (b) of Figure 9 instead presents the full RIMP distribution implied by the two calibrations. In addition, it superimposes the empirical RIMP distribution, the object we present in Figure 1. It is evident that the two calibrations yield drastically different RIMP distributions. Both distributions have a large mass of firm-time observations at low, but distinct from zero, RIMP values. Version A delivers a more concentrated RIMP distribution, which trails off towards larger RIMPs, but lacks a long tail. Version B is bimodal, with a significant mass at RIMP values that are large in absolute value.

Since the RIMP is a general equilibrium object, several economic forces interact in generating the patterns in Figure 9. It is helpful to separate the distributions in Panel (b) into firms for which the collateral constraint binds and those for which it does not.

[14] Focusing on the size-RIMP correlation is not precisely the same as obtaining a coefficient on the interaction term between size and monetary policy surprise in an investment panel regression. Targeting a simpler univariate correlation moment retains the spirit of selecting on firm characteristics *ex ante* while being slightly easier to compute in the model.

Figure 9: RIMP-size relation and RIMP distribution across different calibrations



(a) RIMP-size relationship

(b) Full RIMP distribution

Notes: The RIMP is constructed using data generated from the stationary distribution of the model of firm heterogeneity in Winberry et al. (2025). See Table 6 for further information about the calibrations.

It turns out that in Version A, the differences in the RIMP between these two groups of firms are less pronounced than in Version B, where investment of constrained firms is very responsive to monetary policy surprises. Importantly, these differences in the RIMP distributions can arise while the RIMP-size correlation remains the same. The reason is that the characteristics of constrained and unconstrained firms, in particular in the size dimension, are also different across the two model versions. These multidimensional differences are economically meaningful features of the model, but something that a calibration strategy based on the standard one-covariate approach would ignore.

Our model experiment demonstrates that our empirical findings are not only relevant for understanding firm behavior in the transmission of monetary policy in its own right. Our innovation of estimating the full RIMP distribution also provides critical empirical targets for researchers who develop models of the investment channel of monetary policy. Our calibration exercise makes clear that the RIMP distribution is useful to select among model calibrations that are economically quite different. For instance, in Version B most of the investment responsiveness of monetary policy works through financially constrained firms, which is not the case in Version A.

In addition to choosing among alternative calibrations of the same model, our estimate of the RIMP distribution could be useful for structural models in three additional ways. First, the RIMP distribution can help to distinguish among models with different ingredients. In fact, our empirical results can guide researchers to decide about the right ingredients for their structural macroeconomic models of firm dynamics. For example, non-convex capital

adjustment costs are likely to affect the shape of the RIMP distribution. Lumpy behavior induced by fixed adjustment costs will likely generate large negative RIMPs, especially when monetary shocks make firms switch away from inaction. On the other hand, firms with large investment inaction regions would likely have a RIMP of zero, a feature that is not borne out by the data. Second, our results support models with limited permanent heterogeneity, but meaningful firm-specific time variation in the RIMP, such as standard models in which heterogeneous firms cycle through different financial constraint regimes. The nature of financial constraints or non-convex adjustment costs are likely to affect the persistence of the RIMP, but have so far mostly been disciplined by univariate heterogeneity in the RIMP. Third, our result that observable characteristics explain little of the overall RIMP heterogeneity opens the door to entirely new model mechanisms, potentially supported by newly collected data on firm behavior.

5 Conclusion

We estimate the distribution of firms' investment responses to changes in monetary policy. Our clustering regression approach contrasts the typical approach in the literature, which considers heterogeneity along a small number of firm characteristics, such as firm size or leverage. It can capture multidimensional and unobservable heterogeneity across firms and time. We find that the investment of most firms in most time periods responds only weakly to interest rate changes. Only about 5% of firm-time observations are associated with a strong responsiveness. In those cases, a 25 bp interest rate hike lowers the quarterly growth rate of firm capital by about 1 pp on average. We relate our estimates to observable firm characteristics *ex post*. While characteristics studied in the existing literature play a role, most of the responsiveness of investment corresponds to variation within firms over time that cannot be explained with observable characteristics. We show that our findings can be used to discipline structural models of heterogeneity in firm investment and monetary policy.

References

Abel, Andrew B, "Optimal investment under uncertainty," *The American Economic Review*, 1983, *73* (1), 228–233.

_ **and Janice C Eberly**, "A unified model of investment under uncertainty.," *American Economic Review*, 1994, *84* (5).

Abel, Andrew B. and Janice C. Eberly, "Optimal Investment with Costly Reversibility," *The Review of Economic Studies*, 10 1996, *63* (4), 581–593.

Almuzara, Martín and Víctor Sancibrián, "Micro responses to macro shocks," *FRB of New York Staff Report*, 2024, (1090).

Aruoba, S Borağan and Thomas Drechsel, "The long and variable lags of monetary policy: Evidence from disaggregated price indices," *Journal of Monetary Economics*, 2024, *148*, 103635.

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

Baley, Isaac and Andrés Blanco, "Aggregate dynamics in lumpy economies," *Econometrica*, 2021, *89* (3), 1235–1264.

Bandeira, Miguel, "Disaggregated impulse response functions via the classifier-Lasso," Technical Report 2023.

Bauer, Michael D and Eric T Swanson, "An alternative explanation for the "fed information effect"," *American Economic Review*, 2023, *113* (3), 664–700.

Best, Lea, Benjamin Born, and Manuel Menkhoff, "The impact of interest: Firms' investment sensitivity to interest rates," Technical Report, ECONtribute Discussion Paper 2026.

Bils, Mark, Peter J Klenow, and Cian Ruane, "Misallocation or mismeasurement?," *Journal of Monetary Economics*, 2021, *124*, S39–S56.

Bonhomme, Stéphane and Elena Manresa, "Grouped Patterns of Heterogeneity in Panel Data," *Econometrica*, 2015, *83* (3), 1147–1184.

Caballero, Ricardo J and Eduardo MRA Engel, "Explaining investment dynamics in US manufacturing: a generalized (S, s) approach," *Econometrica*, 1999, *67* (4), 783–826.

Caglio, Cecilia R, R. Matthew Darst, and Sebnem Kalemli-Ozcan, "Collateral Heterogeneity and Monetary Policy Transmission: Evidence from Loans to SMEs and Large Firms," Working Paper 28685, National Bureau of Economic Research 2024.

Cattaneo, Matias D, Richard K Crump, Max H Farrell, and Yingjie Feng, "On binscatter," *American Economic Review*, 2024, *114* (5), 1488–1514.

Cloyne, James, Clodomiro Ferreira, Maren Froemel, and Paolo Surico, "Monetary policy, corporate finance, and investment," *Journal of the European Economic Association*, 2023, *21* (6), 2586–2634.

Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva, "The how and why of household reactions to income shocks," Technical Report, National Bureau of Economic Research 2024.

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

David, Joel M, Raffaella Giacomini, Xiyu Jiao, and Weining Wang, "Causal State-Dependent Local Projections," *arXiv preprint arXiv:2605.05404*, 2026.

Deng, Minjie and Min Fang, "Debt maturity heterogeneity and investment responses to monetary policy," *European Economic Review*, 2022, *144*, 104095.

Fukui, Masao, Niels Joachim Gormsen, and Kilian Huber, "Sticky Discount Rates," Technical Report, National Bureau of Economic Research 2024.

Gertler, Mark and Simon Gilchrist, "Monetary policy, business cycles, and the behavior of small manufacturing firms," *The Quarterly Journal of Economics*, 1994, *109* (2), 309–340.

González, B., G. Nuño, D. Thaler, and S. Albrizio, "Firm Heterogeneity, Capital Misallocation and Optimal Monetary Policy," CEPR Discussion Paper 18695, CEPR, Paris and London 2023.

Gormsen, Niels Joachim and Kilian Huber, "Corporate discount rates," *American Economic Review*, 2025, *115* (6), 2001–2049.

Graham, John R. and Campbell R. Harvey, "The Duke University CFO Outlook Survey: 1996–2020," 2020.

Haltiwanger, John, Ron S Jarmin, and Javier Miranda, "Who creates jobs? Small versus large versus young," *Review of Economics and Statistics*, 2013, *95* (2), 347–361.

Hayashi, Fumio, "Tobin's marginal q and average q: A neoclassical interpretation," *Econometrica: Journal of the Econometric Society*, 1982, pp. 213–224.

Huang, Jiaming, "Group Local Projections," *SSRN Electronic Journal*, 2021.

Jeenas, Priit, "Firm balance sheet liquidity, monetary policy shocks, and investment dynamics," 2023.

Jorgenson, Dale W, "Capital theory and investment behavior," *The American economic review*, 1963, *53* (2), 247–259.

Jungherr, Joachim, Matthias Meier, Timo Reinelt, and Immo Schott, "Corporate debt maturity matters for monetary policy," *International Finance Discussion Paper*, 2024, (1402).

Kaplan, Greg and Giovanni L Violante, "The marginal propensity to consume in heterogeneous agent models," *Annual Review of Economics*, 2022, *14* (1), 747–775.

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

Kiyotaki, Nobuhiro and John Moore, "Credit cycles," *Journal of political economy*, 1997, *105* (2), 211–248.

Koby, Yann and Christian Wolf, "Aggregation in heterogeneous-firm models: Theory and measurement," *Manuscript, July*, 2020.

Krusell, Per, Claire Thürwächter, and Joshua Weiss, "Upwardly Mobile: The Response of Young vs. Old Firms to Monetary Policy," Technical Report, Working Paper 2023.

Lewis, Daniel, Davide Melcangi, and Laura Pilossoph, "Latent Heterogeneity in the Marginal Propensity to Consume," *The Review of Economic Studies*, 02 2026, p. rdaf102.

Melcangi, Davide, "The marginal propensity to hire," *FRB of New York Staff Report*, 2025, (875).

_ **and Silvia Sarpietro**, "Nonlinear firm dynamics," *Working paper*, 2026.

Ottonello, Pablo and Thomas Winberry, "Financial heterogeneity and the investment channel of monetary policy," *Econometrica*, 2020, *88* (6), 2473–2502.

Quandt, Richard E, "A new approach to estimating switching regressions," *Journal of the American statistical association*, 1972, *67* (338), 306–310.

Selgrad, Julia and Kerry Siani, "Monetary policy and investment plans," *Available at SSRN*, 2025.

Winberry, Thomas, Adrien Auclert, Matthew Rognlie, and Ludwig Straub, "New Keynesian Economics with Household and Firm Heterogeneity," Technical Report, National Bureau of Economic Research 2025.

Wroblewski, Caleb, "The interest rate elasticity of investment: micro estimates and macro implications," Technical Report, Mimeo, UC Berkeley 2025.

Zwick, Eric and James Mahon, "Tax policy and heterogeneous investment behavior," *American Economic Review*, 2017, *107* (1), 217–248.

ONLINE APPENDIX TO

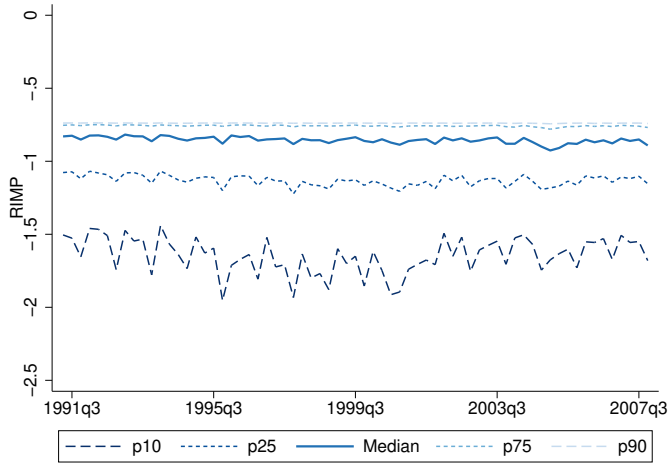
"The Investment Channel of Monetary Policy: Disentangling Firm Heterogeneity"

by Thomas Drechsel, Daniel Lewis, Davide Melcangi and Laura Pilossoph

A Additional results

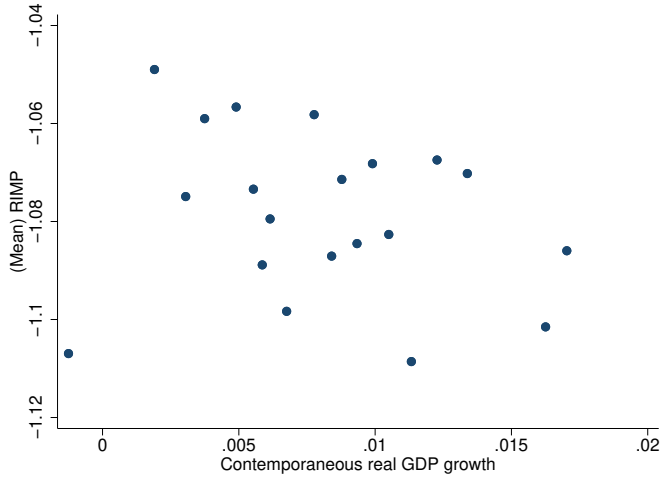
A.1 The RIMP over time and over the business cycle

Figure A.1: Moments of the predicted RIMP distribution over time



Notes: At each quarter, the figure plots the 10th, 25th, 50th, 75th, and 90th percentiles of the cross-sectional distribution of the predicted RIMP across firm-quarter observations.

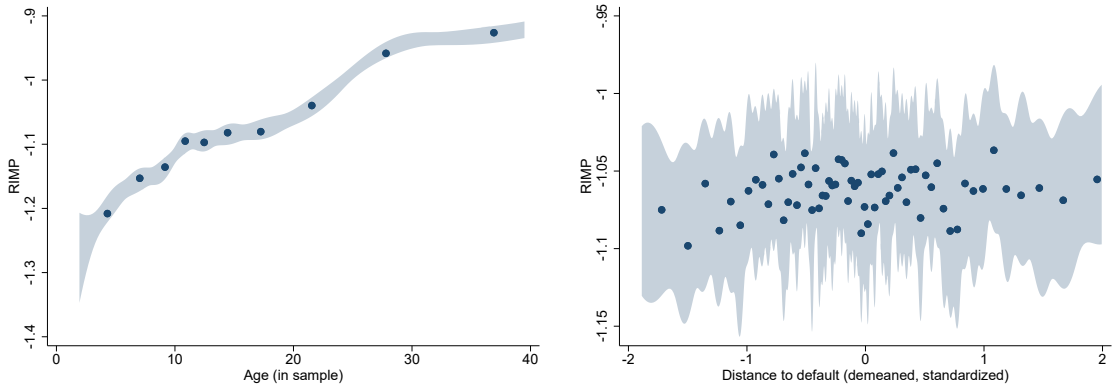
Figure A.2: Average RIMP and real US GDP growth



Notes: The predicted RIMP is averaged across firms for each quarter and then plotted against contemporaneous real GDP growth using a binscatter with 20 quantiles.

A.2 Additional results for the relation with firm observables

Figure A.3: RIMP heterogeneity by firm age and distance to default

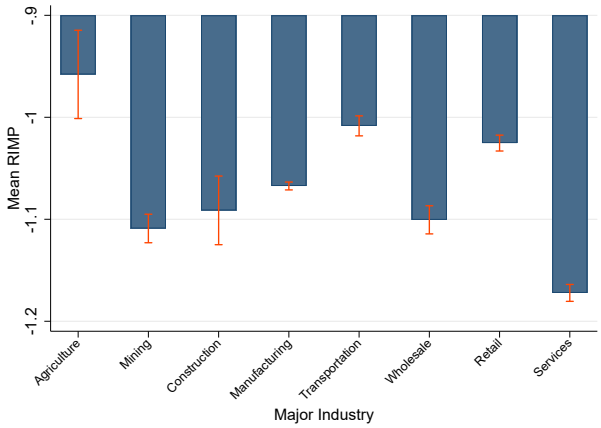


(a) RIMP and age

(b) RIMP and distance to default

Notes: The figure plots Binsreg regressions with 20 quantiles for firm age (defined as years in sample, left panel) and distance to default (computed as in Ottonello and Winberry, 2020, right panel) for our baseline horizon-0 RIMP estimates. Shaded regions report 95% uniform confidence bands for the estimated conditional function.

Figure A.4: RIMP heterogeneity by industry



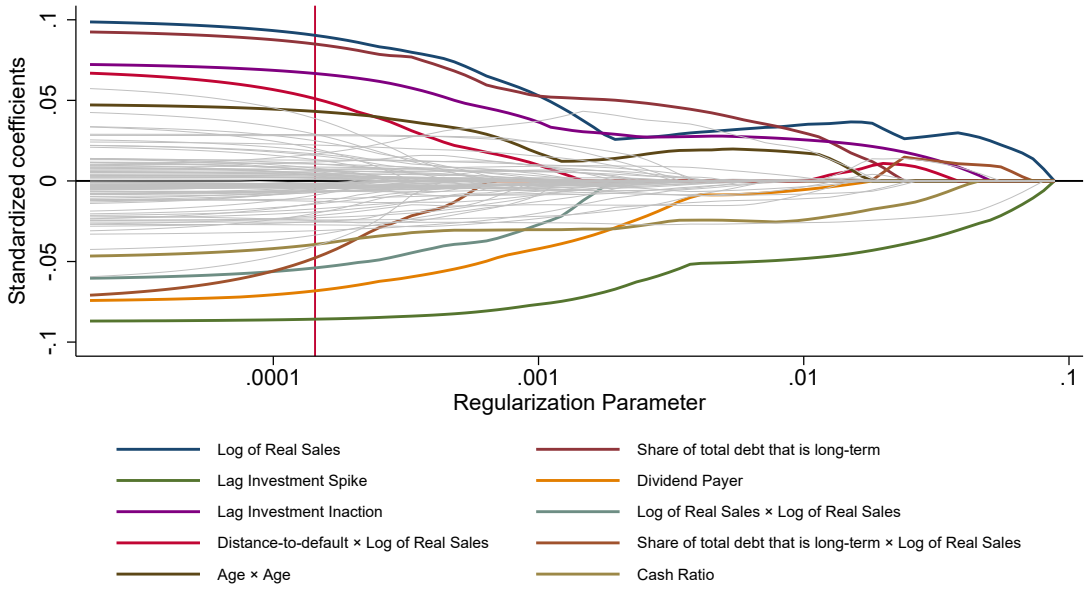
Notes: The bars plot the average predicted RIMP by major industry categories for our baseline horizon-0 RIMP estimates with time-varying group membership. 95% confidence intervals are reported in red.

Table A.1: RIMP and observable characteristics - multivariate analysis (extended version)

	(1)	(2)	(3)
Leverage (demeaned, standardized)	-0.003	-0.003	-0.001
	(0.003)	(0.003)	(0.004)
Distance to default (demeaned, standardized)	0.011***	0.015***	0.004*
	(0.002)	(0.002)	(0.002)
Share of total debt that is long-term	0.046***	0.047***	0.023**
	(0.007)	(0.007)	(0.009)
Cash ratio	-0.264***	-0.261***	-0.128***
	(0.015)	(0.016)	(0.029)
Firm age	0.002***	0.002***	0.000
	(0.000)	(0.000)	(0.001)
Log of real sales (avg. past year)	0.018***	0.023***	0.016**
	(0.001)	(0.002)	(0.007)
Market-to-book ratio	-0.010***	-0.009***	-0.002
	(0.002)	(0.002)	(0.002)
Investment spike	-0.149***	-0.146***	-0.087***
	(0.007)	(0.007)	(0.007)
Investment inaction	0.065***	0.064***	0.030***
	(0.004)	(0.004)	(0.004)
TFP	0.031***	0.011***	0.059***
	(0.003)	(0.004)	(0.012)
Log MRPK	-0.050***	-0.058***	-0.065***
	(0.002)	(0.002)	(0.008)
Quarters since last investment spike	0.003***	0.003***	-0.000
	(0.000)	(0.000)	(0.000)
Observations	90058	90058	89925
R-squared	0.072	0.085	0.205
Adjusted R-squared	0.072	0.080	0.168
Time x Industry FE		✓	
Firm FE			✓

Notes: This table corresponds to Table 3 in the main text, but with additional observables added. These are motivated by the firm dynamics and investment literature. The market-to-book ratio is defined as total liabilities plus the market value of equity, divided by total assets. Investment inaction is a dummy that takes 1 if a growth rate of capital is less than 1% in absolute terms. Investment spike is a dummy that takes 1 if a capital growth rate is greater than 5%. TFP is constructed as in Melcangi and Sarpietro (2026). Our proxy for MRPK is constructed similarly to González et al. (2023), but we use sales instead of value added because our panel does not consistently contain a measure of value added.

Figure A.5: RIMP heterogeneity: LASSO



Notes: The figure plots the evolution of standardized coefficients as the regularization parameter increases for a LASSO regression of our baseline impact predicted RIMP estimates with time-varying group membership on several predictors and their pairwise interactions at different regularization parameters. The red solid vertical line indicates the optimized regularization parameter; in the legend, we list the 10 largest predictors at this regularization level. "Inaction" is defined as having an absolute capital growth rate less than 1%. "Investment spike" is defined as having a capital growth rate higher than 5%. "Dividend payer" is a dummy taking a value of 1 if the firm paid dividends for at least one quarter in the last year. All other variables are defined in the text.

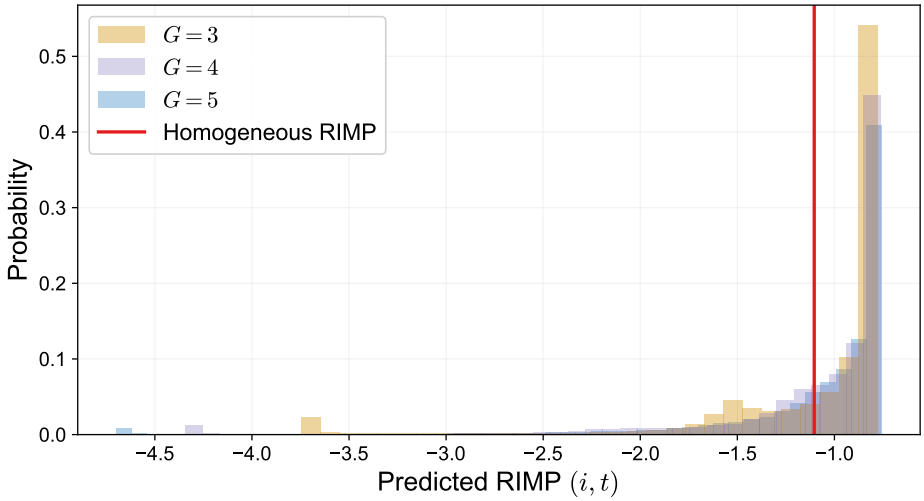
Table A.2: RIMP and observable characteristics - group-specific predictors

	Group 1 $\lambda_1 = -4.21$	Group 2 $\lambda_2 = -1.66$	Group 3 $\lambda_3 = -0.90$	Group 4 $\lambda_4 = -0.51$
Leverage (demeaned, standardized)	0.000	-0.000	0.000	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Distance to default (demeaned, standardized)	-0.000	-0.001*	-0.001	0.002***
	(0.000)	(0.001)	(0.001)	(0.001)
Share of total debt that is long-term	-0.014***	-0.032***	0.005***	0.042***
	(0.001)	(0.002)	(0.002)	(0.002)
Cash ratio	0.038***	0.114***	0.027***	-0.178***
	(0.002)	(0.004)	(0.004)	(0.005)
Firm age	-0.001***	-0.002***	0.000	0.002***
	(0.000)	(0.000)	(0.000)	(0.000)
Log of real sales (avg past year)	-0.004***	-0.007***	0.002***	0.009***
	(0.000)	(0.000)	(0.000)	(0.000)
Pseudo R^2	0.027	0.015	0.000	0.014

Notes: In each column the dependent variable is the predicted RIMP at horizon 0. The reported coefficients are marginal effects from a fractional multinomial logit estimation. The reported pseudo-R^2 is the McFadden pseudo-R^2 for the full fractional multinomial log-likelihood. Robust standard errors are reported in parentheses. *,**, and *** denote significance at the 10, 5, and 1% levels respectively.

A.3 Robustness for choice of number of groups G

Figure A.6: Predicted RIMP distributions for different G



Notes: The plot overlays the predicted posterior-weighted RIMP histogram at the firm-quarter level generated by different choices of the number of groups, G. The displayed homogeneous RIMP is the same as in Figure 1 and the panel LP in Equation (1).

Table A.3: RIMP and observable characteristics - robustness to G

	Number of groups (G)								
	3			4			5		
Leverage (demeaned, standardized)	-0.001	-0.002	0.002	-0.001	-0.002	0.002	-0.001	-0.002	0.002
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Distance to default (demeaned, standardized)	0.004**	0.005**	0.001	0.004**	0.005**	0.000	0.004**	0.005**	0.000
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Share of total debt that is long-term	0.094***	0.096***	0.031***	0.096***	0.098***	0.035***	0.096***	0.098***	0.035***
	(0.006)	(0.006)	(0.008)	(0.006)	(0.006)	(0.008)	(0.006)	(0.006)	(0.009)
Cash ratio	-0.331***	-0.343***	-0.146***	-0.330***	-0.342***	-0.142***	-0.329***	-0.341***	-0.141***
	(0.013)	(0.014)	(0.026)	(0.014)	(0.014)	(0.027)	(0.014)	(0.014)	(0.028)
Firm age	0.004***	0.004***	0.002***	0.004***	0.004***	0.001**	0.004***	0.004***	0.001**
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.001)
Log of real sales (avg past year)	0.022***	0.023***	0.014***	0.022***	0.023***	0.017***	0.023***	0.023***	0.018***
	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.005)	(0.001)	(0.001)	(0.005)
Observations	110395	110379	110307	110395	110379	110307	110395	110379	110307
R-squared	0.043	0.056	0.213	0.041	0.053	0.206	0.039	0.051	0.200
Adjusted R-squared	0.043	0.051	0.179	0.041	0.049	0.172	0.039	0.047	0.165
Time x Industry FE		✓			✓			✓	
Firm FE			✓			✓			✓

Notes: This table repeats the specifications reported in Table 3 for alternative numbers of groups, $G = 3, 4, 5$. Within each value of G, the three columns correspond to the same specifications as in Table 3. The dependent variable is the predicted horizon-0 RIMP. Regressors are lagged by one quarter. Robust standard errors are reported in parentheses. *, **, and *** denote significance at the 10, 5, and 1% levels, respectively.

A.4 Specification with group-specific intercepts

In this appendix, we report the results when we estimate the following model:

$$\widetilde{\Delta \log k_{i,t}} = \sum_{g=1}^{G} d_{it,g} \left[\alpha_g + \lambda_g \widetilde{\epsilon_t^m} \right] + \widetilde{\eta}_{i,t}, \quad \widetilde{\eta}_{i,t} \sim \mathcal{N}\left(0, \sigma_g^2\right). \tag{7}$$

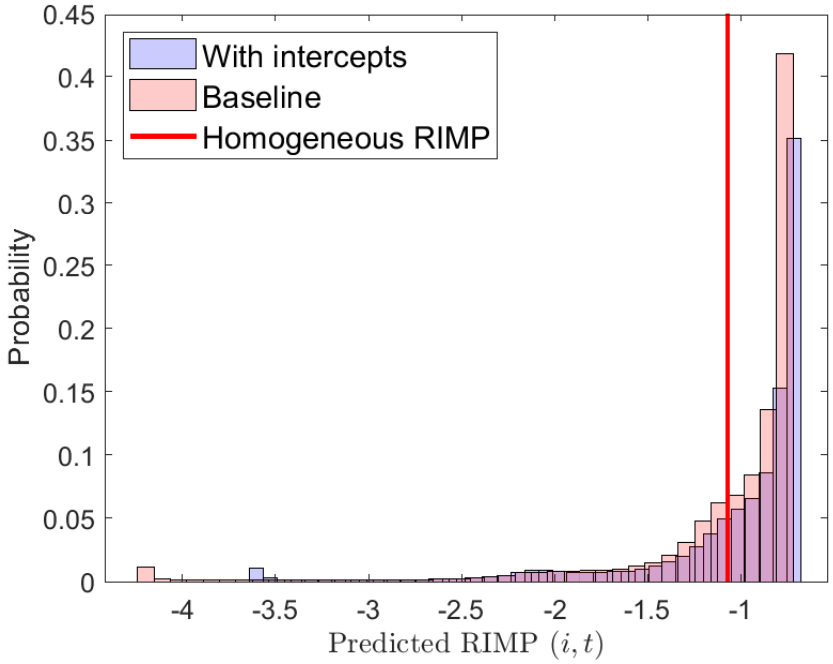
This model differs from our main specification by including group-specific intercepts, α_g. Table A.4 and Figure A.7 present the unconditional and predicted RIMP distributions for specification (7). The estimates are quite close to our baseline results shown in Table 1 and Figure 1 of the main text, respectively. The biggest difference is found in the group of the largest responders: while it still comprises 5% of observations, the RIMP in this robustness exercise is -3.6 while in our baseline it is -4.2. The predicted RIMPs in the two estimation procedures have a correlation of 0.99.

Table A.4: Robustness to group-specific intercepts – RIMP unconditional distribution

	group 1	group 2	group 3	group 4
RIMP (λ_g)	-3.59	-1.84	-0.81	-0.50
Standard Error	(1.62)	(0.34)	(0.11)	(0.06)
Probability of being in group (π_g)	0.05	0.19	0.45	0.31
Standard Error	(0.00)	(0.01)	(0.01)	(0.01)

Notes: Estimates of the horizon-0 RIMP and group membership probabilities in the model including group-specific intercepts, Equation (7). Standard errors are in parentheses, based on the Fisher Information clustered by time.

Figure A.7: Robustness to group-specific intercepts – predicted distribution



Notes: The histogram plots the predicted RIMP at the firm-quarter level, as computed by Equation (5), for the baseline model and the one with group-specific intercepts. The homogeneous RIMP is the same in the two models.

B Omitted variables bias in interacted specifications

In this section, we explain how projecting the distribution of the RIMP on observables may obtain different results from interacted specifications, like those of OW. For clarity, we consider a simplified version of the OW specification (Equation (2)):

$$\Delta \log k_{i,t} = \beta x_{i,t} \epsilon_t^m + u_{i,t}. \tag{8}$$

Without loss of generality, assume that all variables are mean zero. Now suppose that in the true model, there are two relevant determinants of RIMP heterogeneity, $x_{i,t}$ and $y_{i,t}$, with $y_{i,t}$ omitted from Equation (8):

$$\Delta \log k_{i,t} = \theta x_{i,t} \epsilon_t^m + \delta y_{i,t} \epsilon_t^m + v_{i,t}. \tag{9}$$

Let $\xi_{i,t} \equiv x_{i,t}\epsilon_t^m$ and $\zeta_{i,t} = y_{i,t}\epsilon_t^m$. Then the two models can be written as

$$\Delta \log k_{i,t} = \beta \xi_{i,t} + u_{i,t} \qquad \text{(estimated)} \tag{10}$$

$$\Delta \log k_{i,t} = \theta \xi_{i,t} + \delta \zeta_{i,t} + v_{i,t}. \qquad \text{(true)} \tag{11}$$

This is now a case of classic omitted variables bias. Let μ be the population linear projection coefficient of $\zeta_{i,t}$ on $\xi_{i,t}$, that is $\mu = \mathbb{E}\left[\xi_{i,t}^2\right]^{-1} \mathbb{E}\left[\xi_{i,t}\zeta_{i,t}\right]$. Then $\beta = \theta + \mu\delta$. It is thus possible to obtain a negative value for β using the interacted approach, even if θ is positive, if μ and δ have opposite signs. Note that $y_{i,t}$ need not be an observable firm characteristic, but could be a firm-level state variable unavailable to the econometrician, for instance. Moreover, this result also nests misspecification due to non-linearities, since $y_{i,t}$ could be $x_{i,t}^2$, or, more generally, $y_{i,t} = f(x_{i,t})$.

Now consider our estimation strategy. For simplicity, assume that $x_{i,t}$ and $y_{i,t}$ are both binary, so that there are four possible values for the response of investment to ϵ_t^m (namely, $0, \theta, \delta, \theta + \delta$), and assume that $v_{i,t}$ is Gaussian. Then applying GMLR to the equation

$$\Delta \log k_{i,t} = \sum_{g=1}^{4} d_{it,g} \lambda_g \epsilon_t^m + v_{i,t} \tag{12}$$

will recover $\{\lambda_g\}_{g=1}^4 = \{0, \theta, \delta, \theta+\delta\}$, so we obtain the true RIMPs. Note that this result does not require the econometrician to have any knowledge of $y_{i,t}$. If we knew group membership with certainty, projecting these values on $x_{i,t}, y_{i,t}$ would recover exactly θ and δ, the true interaction coefficients. In practice, using the predicted values, $\hat{\lambda}_{i,t}$, introduces measurement error, which will impact the coefficient estimates. However, in this example, if all of the λ_g

responses have the same sign, as is the case for our estimates, the sign of each projection coefficient is preserved, even if the coefficients are biased.

C Model of firm heterogeneity and monetary policy

This appendix provides more details on the model exercise presented in Section 4.2. The structural model we analyze consists of the following blocks:

(i) a continuum of perfectly competitive *heterogeneous wholesale producers*, producing a single homogeneous wholesale good using capital and labor,

(ii) a continuum of monopolistically competitive *retailers* $i \in [0, 1]$, each turning one unit of the wholesale good into one unit of differentiated variety i, setting nominal prices subject to a Rotemberg adjustment cost,

(iii) a *final-good producer*, CES aggregator over varieties with elasticity γ,

(iv) a *representative household* that consumes, supplies labor with linear disutility, owns all firms in the economy, and lends in nominal bonds,

(v) a *monetary authority* setting the nominal rate via a Taylor rule on inflation.

The problem of heterogeneous wholesale producers is taken from Winberry et al. (2025). The firms can borrow subject to a collateral constraint and face a non-negativity constraint on dividends. They exit with an exogenous probability and exiting firms are replaced by newborn firms, with k_0 units of capital and no debt. We describe the producer maximization problem further below.

The remaining blocks closely follow Ottonello and Winberry (2020). The Taylor rule is $\log(1 + r_t^{\text{nom}}) = \log(1/\beta) + \phi_\pi \log \pi_t + \epsilon_t^m$. There is no aggregate uncertainty and we study perfect-foresight transition paths (MIT shocks, ϵ_t^m). The New Keynesian Phillips curve is

$$\pi_t(\pi_t - 1) = \frac{\gamma - 1}{\psi}\left(\frac{\gamma}{\gamma - 1} p_t - 1\right) + \mathbb{E}_t\left[\Lambda_{t,t+1} \frac{Y_{t+1}}{Y_t} \pi_{t+1}(\pi_{t+1} - 1)\right].$$

C.1 Firm problem for wholesale producers

Wholesale producers solve

$$
\begin{aligned}
v(z,k,b) = & \max_{n,k',b',d} d + \mathbb{E}\Lambda\left[(1-\theta)v(z',k',b') + \theta v_{\text{exit}}(z',k',b')\right] \\
& \text{s.t.} \\
& d = pzk^{\alpha}n^{\nu} - wn - b - (k' - (1-\delta)k) - \Phi\left(k,k'\right) + \frac{b'}{1+r}, \\
& d \geq 0, \\
& b' \leq \chi k', \\
v_{\text{exit}}(z,k,b) = & zk^{\alpha}n^{\nu} - wn - b + (1-\delta)k,
\end{aligned}
$$

where z is idiosyncratic productivity, which evolves following a Markov transition, k is capital, b is net debt, and n is labor. Capital adjustment costs are $\Phi\left(k,k'\right) = \frac{\phi}{2}\left((k'-k)/k\right)^2 k$. p is the wholesale price, which is also the marginal cost of retailers. Λ is the household's stochastic discount factor.

C.2 Full calibration table

Table C.1: Alternative calibrations of the structural model

Parameter	Version A	Version B
Capital adjustment cost parameter (ϕ)	0.20	0.05
Collateral constraint parameter (χ)	0.20	0.70
Taylor rule coefficient on inflation (ϕ_π)	1.25	1.25
Elasticity of substitution (γ)	10	10
Price Adjustment Cost (ψ)	90	90
Capital exponent (α)	0.32	0.32
Labor exponent (ν)	0.60	0.60
Depreciation rate (δ)	0.025	0.025
Discount factor (β)	0.99	0.99
Persistence of productivity (ρ_z)	0.94	0.94
SD of productivity innovations (σ_z)	0.085	0.085
Exit probability (θ)	0.02	0.02
Initial Capital (k_0)	0.5	0.5
Initial Debt (b_0)	0	0

Notes: The parameter governing the household's disutility of labor supply is selected in each model to generate a steady state employment rate of 60%, as in Ottonello and Winberry (2020). Idiosyncratic productivity follows the same process as Winberry et al. (2025): $\log z' = \rho_z \log z + \varepsilon$, with $\varepsilon \sim \mathcal{N}\left(0, \sigma_z^2\right)$. The following parameters are set as in Ottonello and Winberry (2020): ϕ_π, γ, ψ, β, b_0. The following parameters are set as in Winberry et al. (2025): α, ν, δ, ρ_z, σ_z, θ, k_0.