The Investment Channel of Monetary Policy: Disentangling Firm Heterogeneity*

Thomas Drechsel
University of Maryland, NBER, CEPR

Daniel Lewis
University College London, CEPR

Davide Melcangi
Federal Reserve Bank of New York

Laura Pilossoph
Duke University, NBER

March 26, 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 monetary policy surprises, using a clustering regression framework. Our new 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 responsiveness. In those cases, a 25 basis point interest rate hike lowers the quarterly growth rate of firm capital by about 1 pp on average. We then correlate our investment sensitivity estimates to observable firm characteristics. While typical characteristics studied in the literature predict sensitivity, we also uncover more novel correlates. In general, the responsiveness of investment cannot easily be explained with a small number of typical firm-level observables. Our rich empirical estimates are useful to discipline 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, Keshav Dogra, Simon Gilchrist, Kilian Huber, Silvia Miranda-Agrippino, Matthias Meier, Manuel Menkhoff (discussant), Pablo Ottonello, and Thomas Winberry for detailed discussions. Conference and seminar participants at the New York Fed, the SED in Copenhagen, 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

One way in which monetary policy affects economic outcomes is by altering firms' investment incentives. Recent research on the investment channel of monetary policy has focused on understanding heterogeneity in the cross section of firms. Typically, researchers first choose a firm-level characteristic of interest, such as firm size or leverage, and then study whether firm-level investment responds more or less strongly to a monetary policy change 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, and potentially several more, play a role – and possibly interact – in explaining the responsiveness of investment to monetary policy.

The goal of this paper is to unpack 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 and instead employ a statistical approach to characterize the full distribution of firms' investment sensitivities to monetary policy. 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 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*. An ex post analysis allows us to investigate the degree to which firm characteristics studied by the literature contribute to this (fully measured) 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.

Our new approach for disentangling firm heterogeneity is a clustering (Gaussian mixture) regression framework, applied to the panel local projection framework 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 optimally determined 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. (2024). 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. However, since longer horizon responses of investment are of particular interest in the broader macroeconomics literature, we also apply our clustering approach to the impulse response functions of investment at the firm level. To obtain estimates at longer horizons, we develop a novel Gaussian mixture estimator for panel local projections, leveraging a bespoke GLS-type transformation.

We present five sets of findings. The first represents our estimated distribution of investment responses to monetary policy across firm-time observations, or RIMP distribution. We find that most firm-time observations correspond to an economically small responsiveness. Notably, our estimates imply that even this low responsiveness estimate is different from zero in a statistical sense, 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 interest rate hike lowers the quarterly growth rate of firm capital by about 1 percentage point (pp) on average. Compared to the literature, our strongest responsiveness estimate represents an economically large effect, although direct comparison can be complicated by the exact specification and setup (see Koby and Wolf (2020) for a detailed discussion).

The second set of findings reveals the sources of variation in the RIMP. About 80% of the variation in the RIMP arises within firms across time, while about 20% occurs across firms at a given point in time. In particular, the state of being associated with the highest RIMP is highly transient. A firm that responds with a 1 pp investment cut to a 25 basis point 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 60% probability. In combination with the first set of findings, the bottom line from these patterns is that the investment of most firms in most time periods responds little to changes in monetary policy, but a few firms in some periods appear to respond very strongly.

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 these firm characteristics indeed play an important role in explaining heterogeneity. One example is that, 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 part of the variation in the responsiveness of investment to changes in monetary policy remains unexplained by observables, even when jointly considering firm characteristics that have been proposed by the literature. A regression with the RIMP on the left-hand side and several observable firm characteristics as well as industry-time fixed effects on the right-hand side leads to a relatively modest R^2 (around 5%). Even firm fixed effects or a LASSO regression that allows for nonlinearities and interactions of a large number of observable characteristics cannot increase the R^2 above 20%. In other words, typical firm-level observables do not appear to be able to explain much of the heterogeneity in investment sensitivities to monetary policy.

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 long horizons, with a sizable responsiveness for some firms being quite persistent. 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 the multidimensional nature of RIMP heterogeneity at longer horizons.

Our rich findings provide a new set of empirical targets for structural models of heterogeneity in firm investment and monetary policy. As an illustration, we confront a specific structural model with some of our most important empirical findings. The recent survey paper on heterogeneity in macroeconomics by Winberry et al. (2025) serves as our starting point. Using their code, settings, calibration, and definition of a monetary policy shock, we generate counterparts of our empirical results in their model. On the one hand, a comparison of the RIMP distributions reveals important similarities. Both our empirical estimates and the model imply a large mass of observations that display little responsiveness and a longer tail of stronger responsiveness. Interestingly, in both cases, the mass of low responsiveness is to the left of zero: every firm does respond, even if by a small magnitude. In addition, smaller firms are more sensitive to monetary policy in the model as well as in the data. On the other hand, there are important differences, especially quantitative, both in the unconditional RIMP distribution and in the joint correlation of the RIMP with firm size and net leverage.

To summarize, the new empirical findings that we uncover in this paper are relevant for understanding firm behavior in the transmission of monetary policy in their own respect. In addition, the illustrative comparison between the empirical evidence and a structural model makes clear how theories of firm heterogeneity and monetary policy can be confronted with and informed by the findings from our econometric methodology.

Contribution 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 known 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 settings. Earlier contributions along those 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.

Second, methodologically we contribute to work that uses clustering regression framework for questions of economic heterogeneity. Lewis et al. (2024), employ a similar approach to study household MPCs. The Gaussian mixture approach has its origin in Quandt (1972). We develop a novel extension to local projections. Alternative approaches to clustering local projection coefficients include Bandeira (2023) (classifier LASSO) and Huang (2021) (K-means); we highlight key differences in the main text. 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 (2024), 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.

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

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

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 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 ex post. We only include a firm fixed effect α_i^h which 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

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

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 low-dimensional vector, and chosen ex ante by the researcher.

2.2 The clustering approach: estimating RIMP heterogeneity

Instead of a specification where a given $x_{i,t-1}$ is selected ex ante, our approach seeks to characterize the entire RIMP distribution. We outline our clustering via GMLR approach by first explaining the general idea and then presenting the formal procedure and estimation, 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 case for $h > 0$ further below.

Main idea. In principle, we want to allow λ^h to be different for each (i, t) observation, so that a firm's RIMP can be possibly unique 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. This means that, 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. Using estimates of π_g and λ_g, 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 discuss alternative ways to achieve the desired dimensionality reduction.

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^h_i, by demeaning both the left and the right hand side of the estimation equation. We denote the resulting orthogonalized variables "$\widetilde{\cdot}$". The model we estimate is

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

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. Since after absorbing firm fixed effects we make no further use of the panel structure, 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

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

where π_g^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 observable data, is

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

where

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

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. (2024) 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 estimated consistently 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*

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

which minimizes the mean squared error of the predicted value of $\Delta\log k_{i,t}$. It amounts to the econometrician's optimal prediction of the change in capital for observation (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 alternatives to our procedure for estimating heterogeneity in the RIMP. One alternative 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. 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. 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 do, we can uncover how the RIMP is *individually* correlated with observables, as well as *joint* correlations. Our ex post approach takes the form

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

where $\tau_{s,t}$ are industry-time fixed effects for industry s at time t, $X_{i,t}$ is an array of time-varying firm-level covariates, 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 in particular examine 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

The total number of groups is an input into the algorithm. $G = 4$ is selected by the Bayesian Information Criterion (BIC); we consider varying G locally as a robustness exercise. The correlation between posterior-weighted RIMPs for $G = 3$ and our baseline $G = 4$ is 0.98; with $G = 5$ it is 0.99. To compute standard errors, we use the Fisher Information. We cluster by time, following the recommendation of Almuzara and Sancibrián (2024).

2.5 Estimation for $h > 0$

Clustering responses estimated via local projections for horizons after impact poses a challenge. To see why, consider the specification in Equation (1) at horizon $h = 1$. The error $\eta^1_{i,t}$ is mechanically serially correlated. It contains the period $t + 1$ structural shocks that also appear in $\eta^1_{i,t+1}$. The problem compounds at longer horizons: at horizon h, $\eta^h_{i,t}$ contains shocks from periods t through $t + h$. In a standard local projection, this affects efficiency, but not consistency, and can be addressed with appropriate standard errors. In our GMLR setting, however, correctly specifying the error process is required to consistently estimate group membership and thus the group-specific RIMP parameters.

The two most natural approaches to handling serially correlated errors in local projections are unappealing in our setting. The first is lag augmentation, following Montiel Olea and Plagborg-Møller (2021). This approach would deviate significantly from the Ottonello and Winberry (2020) benchmark. More importantly, it would complicate the clustering. There is no reason to believe lag coefficients would not vary with group membership, so the scalar structural parameter of interest would be replaced by a large vector of group-specific coefficients. This would both increase the computational burden and impair interpretability, since group membership might be driven by lag coefficients rather than the RIMP itself. Lewis et al. (2024) provide a related discussion. The second approach is to derive and estimate the covariance structure of the errors across time and horizons. However, despite stationarity assumptions, we found the implied covariance matrix to be ill-conditioned in our sample, and applying regularization causes the EM algorithm to break down.

Motivated by Lusompa (2023) and Breitung and Brüggemann (2023), we instead adopt a GLS-type correction that resolves the serial correlation problem by directly controlling for the cumulated shocks that generate it. The basic intuition is straightforward. At $h = 1$, the residual $\eta^1_{i,t}$ contains the period $t+1$ structural shocks, which also appear in the $h = 0$ residual one period later, $\eta^0_{i,t+1}$. If we can include a vector of lagged controls that span shocks from period $t-1$ and earlier, we can then purge the remaining serial correlation from the error term by subtracting $\eta^0_{i,t+1}$, the future residual, and $\lambda^0_g \widetilde{\epsilon^m_{t+1}}$, the impact of future monetary policy shocks, from the dependent variable $\widetilde{\Delta^{+1} \log} k_{i,t}$, so they depend only on period-t shocks. This delivers a Gaussian vector of residuals at time t that are serially uncorrelated, making GMLR tractable. This argument extends recursively to longer horizons. We provide a formal development in Appendix B.

One important drawback of this approach is that group membership must be time-invariant. Our main $h = 0$ results allow group membership to vary across time, but the recursive GLS transformation requires that group membership be fixed within the transformation window. As we discuss below, this assumption causes only modest attenuation of the estimated RIMP distribution relative to the baseline, and our substantive conclusions are unaffected.

3 Results

We organize our results into several parts. We begin by presenting the unconditional and predicted RIMP distribution. Then, we examine sources of RIMP variation, including its observable drivers, by linking our estimates to firm-level observables. We also 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 formally represented by Equation (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 1pp. Interestingly, while this coefficient is small, it is significantly different from zero. In

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

	group 1	group 2	group 3	group 4
RIMP (λ_g)	-4.22	-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.

other words, there is no group that corresponds to a 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 tightening of 25 basis points, 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 contemporaneous ($h = 0$) investment of most firms in most time periods does not respond strongly 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. 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 local mode at the maximum responsiveness of around -4.2.

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

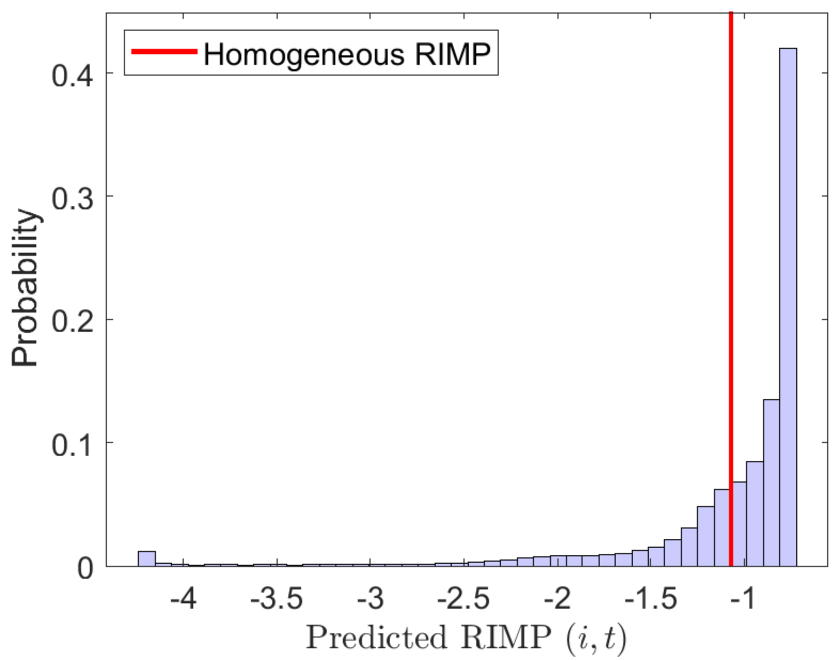

Notes: The histogram plots the predicted RIMP at the firm-quarter level, as computed by Equation (5). Baseline sample of 157,101 firm-quarter observations.

3.2 RIMP variation across firms vs. within firms

We begin by investigating the degree to which RIMP heterogeneity arises across firms as opposed to within firms over time. First, we conduct a simple variance decomposition of the predicted RIMP. We find that about 80% of the variation is within firms and across time. Essentially, responsiveness to monetary policy is a feature that changes over time rather than a fixed firm characteristic. We return to this point in Section 3.3.

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

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)		-4.22	-1.66	-0.90	-0.51	Total
λ_1	-4.22	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	0.44	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, that corresponding to the group with the highest posterior probability for observation j.

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 0.5pp 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 high RIMP status is very much transient. A firm that cuts investment by 1 pp in response to a 25 basis point 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 little, on the other hand, will also respond little in the following period with a 59% probability.

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 most classic explanations for differences in the responsiveness to monetary policy is firm size, following for example Gertler and Gilchrist (1994). Figure 2a 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. 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 2a 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 2b, 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).[1]

Figure 2: RIMP heterogeneity by firm size and age

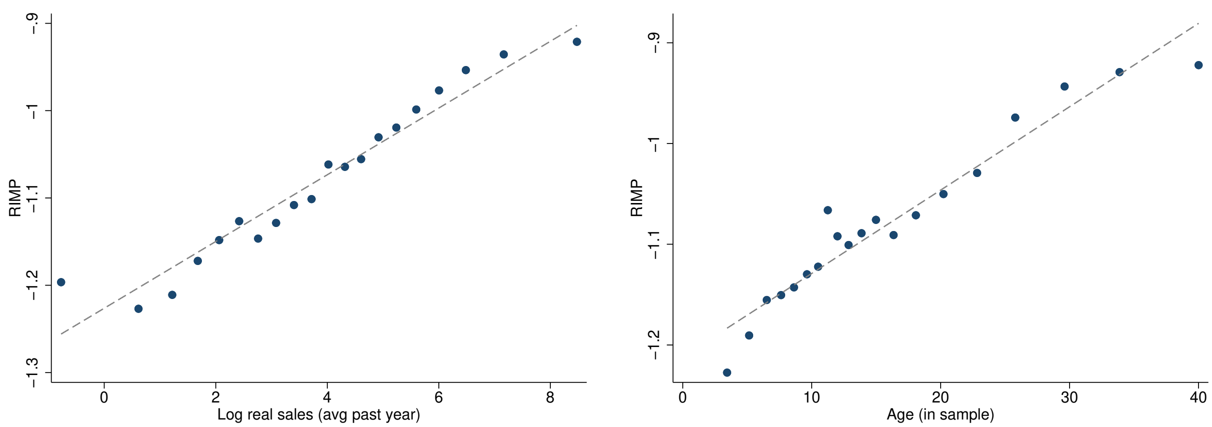

(a) RIMP and size

(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 indicated 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 3 presents 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 3a 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 a weak and unclear relationship, whereas OW find that less risky firms (i.e. those who are distant to default) are more responsive to monetary policy. We discuss possible reasons behind this disconnect below.

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

In Figure 3c, 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 3d that firms with shorter debt maturity are more sensitive to monetary policy, confirming results by Jungherr et al. (2024).[2]

In general, there are several reasons that the interaction term approach followed by the previous literature may imply a different relationship than the correlation between the RIMP and a given firm characteristic.[3] First, the "true" heterogeneity in the RIMP could be highly nonlinear or multi-dimensional, and thus driven by factors other than those captured by a selected observable characteristic such as distance to default. A simple linear interaction ignores these other factors, potentially leading to omitted variable bias. We provide an example in Appendix A. 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. More importantly, a substantial portion of heterogeneity in the RIMP could be driven by *unobserved* factors; these will always be omitted, unless an approach like ours is followed.

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 Figure 2 and 3. We have documented that the RIMP ranges from -0.5 to -4.2. However, in Figure 2a, 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 RIMP difference between having only long-term debt and no long-term debt is about 0.15pp, less than half of the inter-quartile range of RIMPs that we estimate.

Even when considered jointly, observable characteristics 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 the following two columns we also control for industry-time fixed effects (column (2)) and firm fixed effects (column (3)). When tested jointly, the statistical association between firm-level predictors and the RIMP remains. However, the R^2, which can be interpreted as a direct measure of the share of the overall RIMP heterogeneity explained by predictors, is only around 5%. Adding industry and time

[2] 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.

[3] 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 3: RIMP and different firm characteristics

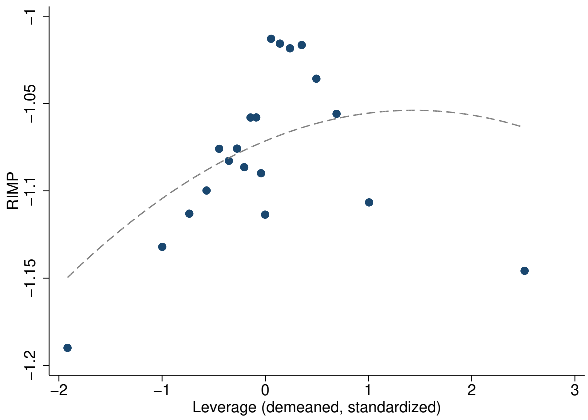

(a) RIMP and leverage

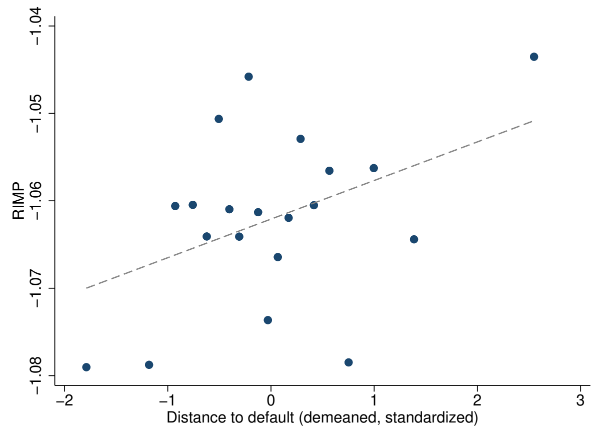

(b) RIMP and distance to default

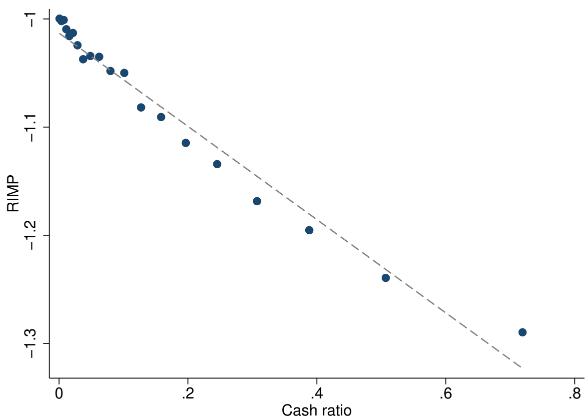

(c) RIMP and cash ratio

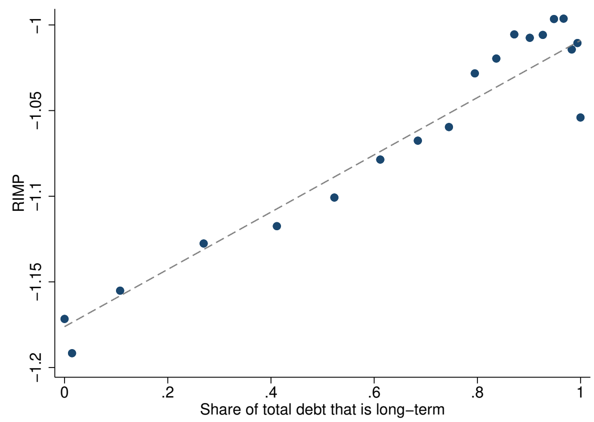

(d) RIMP and debt maturity

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 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.

heterogeneity does not substantially increase the R^2. 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 C.3). Nevertheless, the extent of this variation is economically and statistically small. Even a LASSO regression that allows

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	-0.331***	-0.343***	-0.143***
	(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***	0.016***
	(0.001)	(0.001)	(0.005)
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. Robust standard errors are reported in parentheses. *,**, and *** denote significance at the 10, 5, and 1% levels respectively.

for nonlinearities and interactions in observable characteristics fails to increase the R^2 above 9%. 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 C.3. In addition, we find that firms that were inactive in the quarter before the shock – i.e. whose capital growth rate was less than 1% in absolute terms – as well as those who experienced an investment spike have a RIMP closer to zero, whereas firms that experienced an investment spike before the monetary surprise are more investment-elastic. Finally, adding firm fixed effects (column 3 of Table 3) increases the R^2 to 17%.

3.3.2 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

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 natural 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.[4] We document this in Figure 4a, 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.

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 "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 4b, 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.[5] 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).

We find that the RIMP is significantly correlated with self-reported expectations. Precisely, the survey asks CFOs to rate their optimism about the financial prospects of their own firm on a 0–100 scale. As shown in Figure 5, firms with less optimistic outlooks

[4] 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.

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

Figure 4: RIMP heterogeneity by discount rate wedges and stickiness

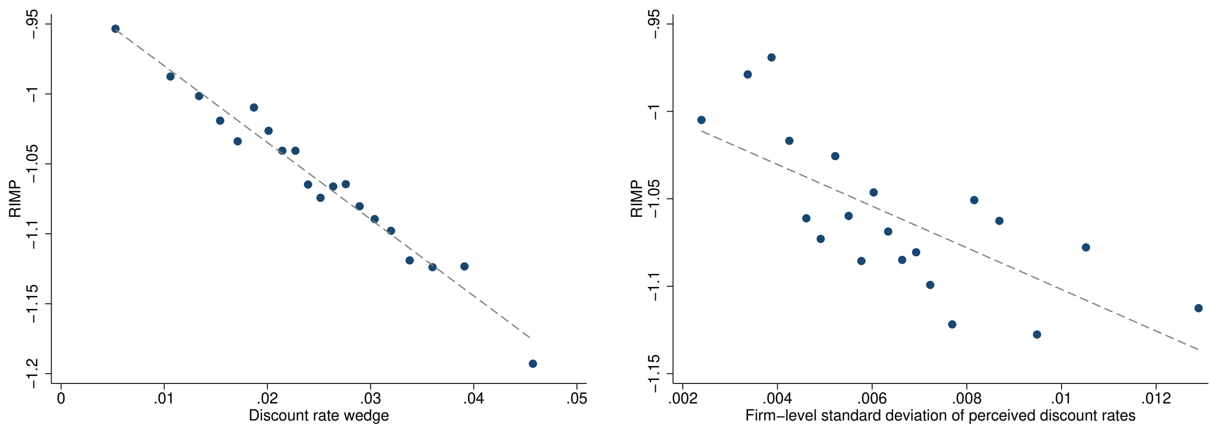

(a) RIMP and discount rate wedges (b) RIMP and (inverse) discount rate stickiness

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 comes fromGormsen and Huber (2025), at quarterly frequency since 2002, yield a sample of 34,952 observations. Dashed lines indicated 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.

about their own prospects are more responsive to monetary policy. The relationship remains statistically significant when we control for other firm-level predictors.[6]

3.3.3 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 6a 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 6b, we focus on two key state variables in heterogeneous firm models: size (defined as lagged log capital, as in the canonical model) and the net leverage ratio (leverage ratio minus cash ratio, as in the model). 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.[7] We discuss in Section 4 how our findings offer a new set of empirical targets that can

[6] 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, but it is statistically insignificant.

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

Figure 5: RIMP heterogeneity by CFO optimism

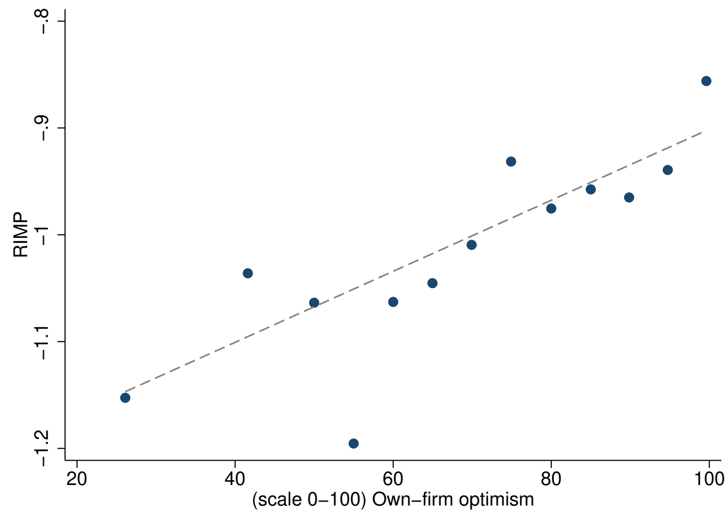

Notes: The figure presents binscatter plots based on 13 quantiles. The predicted RIMP is plotted against a self-reported measure of optimism towards their own firm, on a scale between 0 and 100. The data comes from a vintage of the CFO Business Outlook Survey and has been merged with Compustat using permno anonymized identifier. The data is 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.

directly discipline heterogeneous firm models.

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 C.1. 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 group 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. However, monetary policy is often thought to operate with long and

Figure 6: RIMP multidimensionality

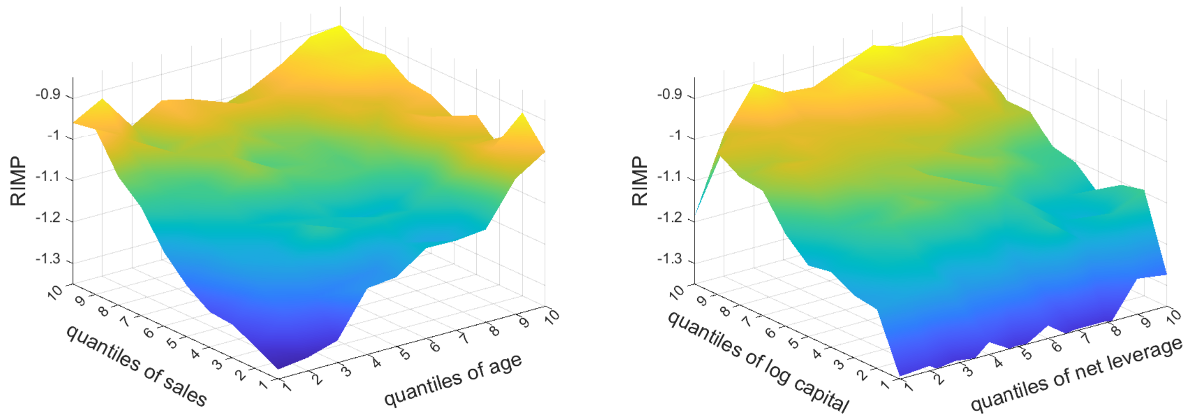

(a) RIMP by size and age

(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 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).

variable lags. It is therefore natural to study the responsiveness of investment to changes in monetary policy at longer horizons.

We approach this question in several steps. First, as discussed in Section 2.5, identification at longer horizons requires the assumption that a firm's group membership does not change over time. As a preliminary step, we therefore re-estimate the RIMP distribution at $h = 0$ imposing this time-invariance assumption. As we show in Appendix C.2, the distribution that we retrieve is similar to our baseline insofar as the largest group of firms has a small, but non-zero, RIMP, with a modest mass of more investment-elastic firms. However, there are two important differences. First, the RIMPs are more concentrated than in our baseline. In particular, the most investment-elastic firms cut capital growth by 2 pp in response to a 1 pp monetary tightening; this semi-elasticity is half what we retrieve in our baseline. This result is consistent with time-invariant group membership pooling time-varying RIMPs within a firm and thus attenuating the true extent of heterogeneity. Indeed, as in Table 2 above, larger investment elasticity is an exceptionally transient phenomenon. Second, the distribution is more discrete, as the posterior weights are largely bimodal. This result is also connected to the time-invariant estimation structure: as the entire sample is now used to classify a firm based on its RIMP, the econometrician's assignment uncertainty falls (i.e. the posterior weights are close to either 0 or 1).[8]

Despite these differences, our results with this approach remains similar to our baseline results even at the firm level. Averaging our baseline (time-varying) predicted RIMP from the previous section at the firm level (i.e. averaging within each firm over time), we find a large positive correlation (0.81) with the time-invariant RIMPs estimated here. As a result, our conclusions on what drives the RIMP are broadly unaffected, as we show in Appendix C.2.

This exercise also addresses potential concerns that the estimated time-varying RIMP may simply capture firms that experience transitory shocks for unrelated reasons in certain periods. If this were the case, when fixing group membership over time those shocks should be subsumed in the error term. The estimated distribution in the RIMP should then be attenuated, and indeed homogeneous if this explanation accounted for all variation in the RIMP. Because the RIMP distribution and its properties do not fundamentally change, we conclude that in our baseline we are identifying heterogeneity in responsiveness to monetary policy, rather than the effects of unobserved idiosyncratic shocks.

As a second exercise, we estimate the RIMP at longer horizons as presented in Section 2.5. For each group, we estimate a RIMP impulse response $\lambda_{g,h}$, $h = 0, \ldots, H-1$. Using estimated posterior weights, we can also construct the firm-specific RIMP impulse response. Put differently, we construct a horizon-specific distribution of the predicted RIMP that is equivalent to what we report in Figure 1. To visualize the results, at each horizon we compute percentiles of this distribution and show them in Figure 7, together with the homogeneous RIMP estimated by imposing $G = 1$.

We find sizable heterogeneity even at longer horizons, and the right skewness of the RIMP distribution appears quite persistent. The most negative RIMP remains sizable, although not quite as large – in absolute terms – as in our baseline results; as discussed earlier, this is the result of time-invariant group membership pooling different RIMPs together.[9] Two years after the monetary policy surprise, some firms do not show almost any investment effect, although the bulk of the RIMPs remain substantially negative.

We then project the predicted RIMP at different horizons on firm characteristics. The results are similar to what we showed before, with small, young, and liquid firms being the most elastic. When we project the RIMP at different horizons on firm predictors jointly, our

[8]Note that in this setting, the BIC exhibits less clear kinks than in our baseline. Hence, we choose the number of groups G using a bootstrap procedure that we describe in Appendix C.2. In the same Appendix, we show how considering a higher G delivers more dispersed and less discrete RIMPs, but does not fundamentally change the conclusions of this section.

[9]As discussed earlier, allowing for more groups slightly increases the range of RIMPs and decreases the discreteness of the distribution, as we show in Appendix C.3. Nevertheless, the shape of the heterogeneous impulse responses is not fundamentally different.

Figure 7: Predicted RIMP: impulse responses

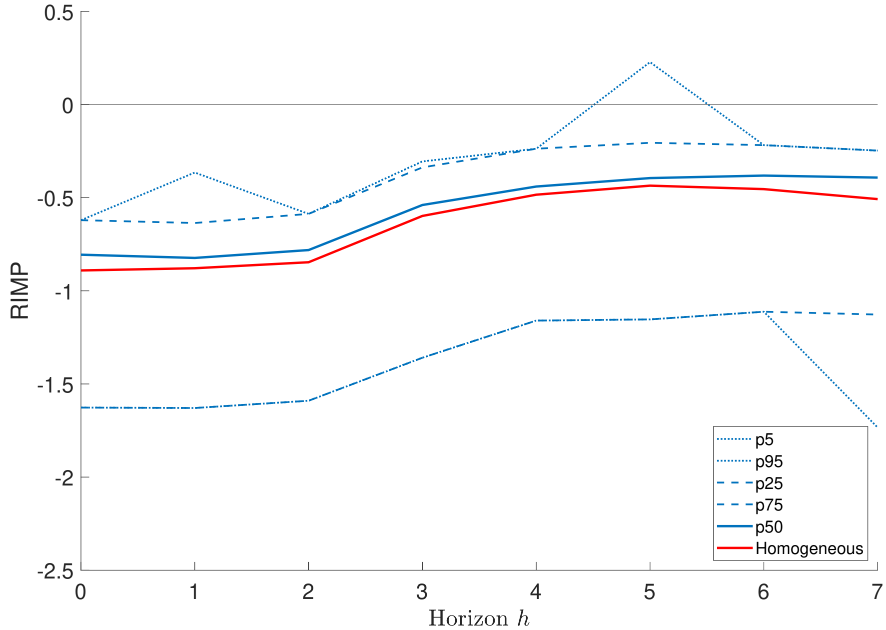

Notes: The red line ("homogeneous") plots the point estimates imposing $G = 1$. For the blue lines, we estimate the GLS-GMLR model described in Appendix C.3 with $G = 4$, compute the predicted RIMP ($\hat{\lambda}_{i,h} = \sum_{g=1}^{G} \gamma_{i,g,h}\lambda_{g,h}$) for each horizon, and then plot the quantiles (5th, 25th, 50th, 75th, and 95th percentiles) of $\hat{\lambda}_{i,h}$ across firms at each horizon. Sample size is restricted to firms with a non-missing dependent variable for all horizons.

conclusions are little affected.[10] The correlations are quantitatively very similar and slightly stronger, whereas the adjusted R^2 is little changed. One difference is that debt maturity is no longer a statistically significant predictor.

4 Implications for structural models

In this section, we discuss the implications of our empirical results for structural models of firm heterogeneity and monetary policy. Our discussion begins by taking stock of our findings, makes broad observations, and presents an explicit comparison between a selection of our findings and their counterparts in a general model from the literature.

4.1 Taking stock of empirical results

The previous section documents a variety of rich empirical patterns uncovered with our new empirical approach. In broad strokes, our findings can be summarized as follows. First, the RIMP distribution is characterized by a mass of firms whose investment responds little to

[10]Given that in this exercise group membership is time-invariant, we cannot add firm fixed effects to the regression.

Table 4: RIMP and observable characteristics - multivariate analysis at longer horizons

	Horizon		
	0	3	7
Leverage (demeaned, standardized)	0.014***	0.013***	0.014***
	(0.002)	(0.002)	(0.002)
Distance to default (demeaned, standardized)	0.001	0.001	0.001
	(0.001)	(0.001)	(0.001)
Share of total debt that is long-term	0.005	0.005	0.006
	(0.004)	(0.004)	(0.004)
Cash ratio	-0.159***	-0.150***	-0.159***
	(0.008)	(0.008)	(0.009)
Firm age	0.003***	0.003***	0.003***
	(0.000)	(0.000)	(0.000)
Log of real sales (avg past year)	0.029***	0.025***	0.032***
	(0.001)	(0.001)	(0.001)
Adjusted R^2	0.069	0.054	0.075

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

monetary policy and a long tail of large responses. Second, most of the variation in the RIMP is within firms rather than across firms. Third, while key firm characteristics studied in the literature are correlated with the RIMP, it is not easy to explain the RIMP with just one single characteristic. Therefore, considering the multidimensional nature of the RIMP is important.

4.2 General considerations around structural models

Our empirical results offer a new set of targets to discipline models with firm heterogeneity and monetary policy. In a burgeoning 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 Kaplan and Violante (2022), Auclert (2019). We hope that our results will enable a similar practice for the RIMP. Many prevalent heterogeneous firm models are quite different from each other, as they are tailored towards the particular firm observable that each paper chooses to focus on, such as distance to default, liquidity, or debt maturity. Nevertheless, in each of these models a RIMP distribution can be computed and multidimensional correlations with the RIMP can be tested. In the next section, we show how to do this in an illustrative example.

There are, however, some important conceptual differences between the MPC and the RIMP that we wish to stress. First, the RIMP is specific to monetary policy. Second, the RIMP is a general equilibrium object. As such, the RIMP is helpful not only for disciplining the features of the firm problem, but also how they interact with equilibrium outcomes and the rest of the economy. In other words, the RIMP is a direct general equilibrium target for monetary policy effects on firm investment, while 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, for instance, Winberry et al. (2025) and Melcangi (2025) for discussions on sensitivities to cash flow shocks, and Zwick and Mahon (2017) and Koby and Wolf (2020) for partial equilibrium investment-interest elasticities. However, empirical evidence on these firm-level partial-equilibrium sensitivities is in scarcer supply than for households and MPCs, making this approach more challenging: the RIMP offers an alternative that is backed by a well-established empirical literature.

A corollary of this discussion is that the RIMP results can also be potentially related to several generations and branches of macroeconomic models in which firm investment is crucial, but which do not explicitly focus on the responsiveness of investment to 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 (1993) and Baley and Blanco (2024), or models of financial frictions, such as Kiyotaki and Moore (1997) and Khan and Thomas (2013). In order to do so effectively, however, it is paramount to take a stance not only on the structure of the firm problem (e.g., financial frictions, adjustment costs), but also on equilibrium (market clearing) conditions, nominal rigidities, and how monetary policy is conducted. The illustrative analysis in the next section does not fully incorporate that: it should therefore be seen as a first step towards mapping the RIMP to structural models.

4.3 Explicit comparison to a structural model

Keeping these considerations in mind, we illustrate how our findings can be useful to modelers. As a case in point, we consider the recent survey paper on heterogeneity in macroeconomics by Winberry et al. (2025) and use it to conduct an explicit comparison with our results. The heterogeneous firm model in Winberry et al. (2025) 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 obtain a RIMP distribution using the code from Winberry et al. (2025), the same settings, same calibration, and same definition of a monetary policy shock as in their paper, which is a one-off upward shift in the real interest rate.[11] As an important caveat, however, we point out that the model-generated RIMPs we present are computed in partial equilibrium (i.e. with a fixed path of endogenous prices) whereas our empirical evidence should be understood as informing general equilibrium objects; see the general discussion above and some concluding remarks later.[12]

Figure 8: Model vs data: the RIMP distribution

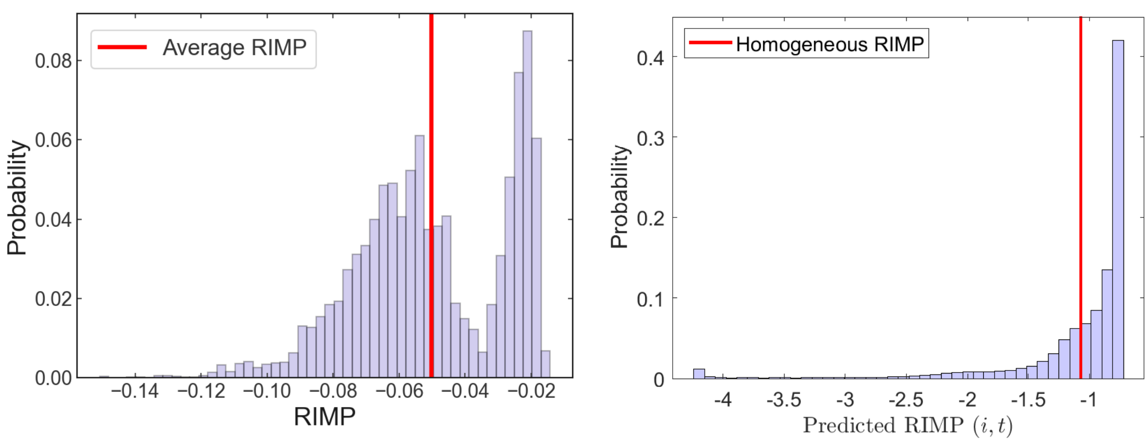

(a) Model (b) Data

Notes: The right panel is the distribution of predicted RIMP in the data (same as Figure 1). The left panel is the model counterpart of Figure 8b. Here, the RIMP is constructed using data generated from the stationary distribution of the model of firm heterogeneity in Winberry et al. (2025).

The first, natural, way to compare model and data is to plot the RIMP distribution itself, which we do in Figure 8. This comparison between the empirical pattern and the model shows several similarities. Both our estimates and the model imply a large mass of

[11] Our empirical results are based on surprises to the nominal interest rate. Apart from this discrepancy between real and nominal rates, we define the RIMP as in our empirical analysis: the change in capital growth rate induced by a 1pp hike in the interest rate.

[12] We thank the authors for making their codes available to us.

Figure 9: Model vs data: the RIMP multidimensionality

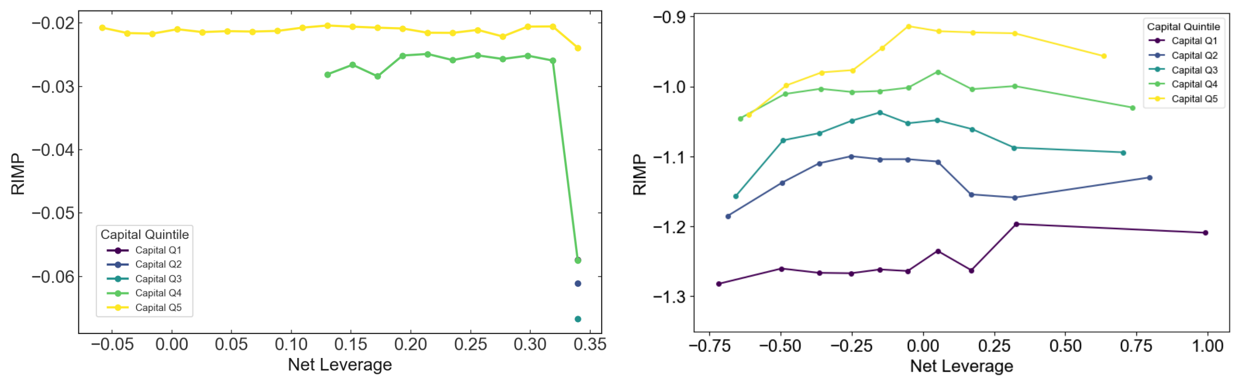

(a) Model

(b) Data

Notes: In the left panel, we use data generated from the Winberry et al. (2025) model of firm heterogeneity. We classify firms, drawn from the stationary distribution, in quantiles of capital (different dotted lines) and net leverage ratio (x-axis). For each cell, we plot the average RIMP (y-axis). Bunching (concentration) of the stationary distribution implies that some slices have fewer dots. In the right panel we do the same in the data, using the RIMP as estimated in Section 3.1. As such, Figure 9b can be seen as an alternative, two-dimensional, representation of Figure 6b.

observations that display little responsiveness and a longer tail of stronger responsiveness. Interestingly, in both cases, the mass of low responsiveness is not right at zero, so every firm does respond, even if by a small magnitude. Overall, the general shape of the RIMP distribution is not dramatically different between data and model.

There are also important differences between our empirical results and the model. For example, the model distribution is bimodal, whereas the empirical distribution declines essentially monotonically from right to left. Importantly, the magnitudes are quite different quantitatively. Our empirical RIMP estimates range from -0.51 to -4.22 while those in the model only range from roughly -0.01 to -0.17. One reason for this divergence may be that the model RIMP is computed in partial equilibrium. More generally, the model results of course depend on the specific calibration.

In Figure 9 we plot the model's (and data's) RIMP by firm size and net leverage. Conceptually, this serves as the structural model counterpart of Figure 6b. Visually, however, we present the results differently because the model-implied stationary distribution is very concentrated at the borrowing constraint (i.e. many firms have the highest possible leverage), making a surface 3D figure harder to read. The different curves in Figure 9a and Figure 9b represent quintiles of firm size (i.e. capital). In the data, the curves shift up monotonically as firms get bigger: smaller firms are more investment-elastic (more negative RIMP). A similar pattern is predicted by the model, albeit much weaker and non-monotonic. Along the horizontal axis, we show how the RIMP changes with net leverage. The model predicts little variation along this dimension, with a more negative RIMP only for the most leveraged firms. In the data this correlation is also weaker than for size. Among net savers (negative net leverage), the empirical RIMP is larger for more liquid firms, especially among larger firms: very few of these firms populate the model-generated stationary distribution. Among net borrowers, the empirical relationship between the RIMP and leverage is less clear, and thus it does not appear substantially different from that implied by the model.

Taken together, this exercise illustrates how structural models can be confronted with our new empirical evidence. Even though we simply follow the calibration of Winberry et al. (2025), some features already align qualitatively between empirical evidence and model. Nevertheless, this analysis should be seen as only illustrative and, importantly, as a first step towards mapping the RIMP to structural models. First, different calibrations will affect RIMP heterogeneity. Second, different adjustment costs (e.g., non-convex, or random as in Winberry, 2021), different financial frictions (e.g., earnings-based constraints as in Drechsel, 2023, or default decisions as in Ottonello and Winberry, 2020), and several other modeling assumptions typically made in firm dynamics models will also affect the RIMP. Third, how general equilibrium forces are modeled, as well as assumptions on monetary policy conduct, will also play a crucial role, as we discussed earlier.

Even under these caveats, this analysis makes clear that multidimensional drivers of the RIMP can easily be studied in structural models, demonstrating the usefulness of our multivariate evidence. While this illustrative comparison focuses on size and leverage as traditional drivers, a promising avenue would be to bring our empirical evidence on more novel predictors, for instance sticky discount rates (Fukui et al., 2024), towards a direct comparison with explicit structural models.

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, and can capture multidimensional and unobservable heterogeneity across firms and time. We find that the investment of most firms in most time periods responds little to monetary policy. Only about 5% of firm-time observations are associated with a strong responsiveness. In those cases, a 25 basis point interest rate hike lowers the quarterly growth rate of firm capital by about 1 pp on average. We relate our investment sensitivity 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 draw conclusions from our findings for 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 Andres Blanco, "The macroeconomics of irreversibility," *The Review of Economic Studies (forthcoming)*, 2024.

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.

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

Breitung, Jörg and Ralf Brüggemann, "Projection Estimators for Structural Impulse Responses," *Oxford Bulletin of Economics and Statistics*, 2023, *85* (6), 1320–1340.

Caballero, Ricardo J and Eduardo MRA Engel, "Microeconomic adjustment hazards and aggregate dynamics," *The Quarterly Journal of Economics*, 1993, *108* (2), 359–383.

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.

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

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

Drechsel, Thomas, "Earnings-based borrowing constraints and macroeconomic fluctuations," *American Economic Journal: Macroeconomics*, 2023, *15* (2), 1–34.

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.

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.

Herbst, Edward P. and Benjamin K. Johannsen, "Bias in local projections," *Journal of Econometrics*, 2024, *240* (1), 105655.

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," Technical Report, National Bureau of Economic Research 2024.

Lusompa, Amaze, "Local projections, autocorrelation, and efficiency," *Quantitative Economics*, 2023, *14* (4), 1199–1220.

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

Olea, José Luis Montiel and Mikkel Plagborg-Møller, "Local Projection Inference Is Simpler and More Robust Than You Think," *Econometrica*, 2021, *89* (4), 1789–1823.

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

Plagborg-Møller, Mikkel and Christian K. Wolf, "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 2021, *89* (2), 955–980.

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, "Lumpy investment, business cycles, and stimulus policy," *American Economic Review*, 2021, *111* (1), 364–396.

_ , Adrien Auclert, Matthew Rognlie, and Ludwig Straub, "New Keynesian Economics with Household and Firm Heterogeneity," Technical Report, National Bureau of Economic Research 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 Omitted variables bias in interacted specifications

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

$$\Delta \log k_{i,t} = \beta x_{i,t}\epsilon_t^m + u_{i,t}. \tag{7}$$

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 (7):

$$\Delta \log k_{i,t} = \theta x_{i,t}\epsilon_t^m + \delta y_{i,t}\epsilon_t^m + v_{i,t}. \tag{8}$$

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

$$\Delta \log k_{i,t} = \beta\xi_{i,t} + u_{i,t} \qquad \text{(estimated)} \tag{9}$$

$$\Delta \log k_{i,t} = \theta\xi_{i,t} + \delta\zeta_{i,t} + v_{i,t}. \qquad \text{(true)} \tag{10}$$

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, $y_{i,t} = f(x_{i,t}) - x_{i,t}$, where $f_{i,t}$ is some functional form such that

$$\Delta \log k_{i,t} = \delta f(x_{i,t})\epsilon_t^m + \nu_{i,t} \tag{11}$$

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 recover 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.

B RIMP Estimation for $h > 0$

Consider the model

$$\log k_t = \Theta(L)\nu_t + B(L)x_t, \tag{13}$$

where $\log k_t$ is the log of the capital stock, which is follows a vector moving average, see for instance the impulse-propagation paradigm in Plagborg-Møller and Wolf (2021). ν_t is a Gaussian vector of nuisance structural shocks, while x_t is the observed shock of interest, the monetary policy surprise. We suppress the dependence of all parameters on group, g, for compactness, as well as the i subscripts on all observations, but the expressions can be taken as conditional on group membership, and to hold in the panel.

The impact change in $\log k_t$ can be written as

$$\Delta \log k_t \equiv \log k_t - \log k_{t-1} = B_0 x_t + \Theta_0 \nu_t + \sum_{l=1}^{\infty} (\Theta_l - \Theta_{l-1})\nu_{t-l} + (B_l - B_{l-1})x_{t-l}. \tag{14}$$

We assume that there exists some vector of controls, W_{t-1}, that spans the past shocks corresponding to non-zero coefficients in Equation (14), so we can write

$$\Delta \log k_t = B_0 x_t + \Theta_0 \nu_t + \Pi W_{t-1}, \tag{15}$$

where B_k and Θ_k are the horizon k coefficients from the lag polynomials above. In practice, our vector of controls, W_{t-1}, contains 4 lags of real GDP growth, the unemployment rate, and CPI inflation as in Ottonello and Winberry (2020), and we assume that Π is group-invariant, so we can partial it out ex ante, following a Frisch-Waugh-Lovell argument. This leaves

$$\Delta \log k_t = B_0 x_t + u_t^0, \tag{16}$$

where $u_t^0 \equiv \Theta_0 \nu_t$. We can similarly write the horizon 1 response as

$$\Delta \log k_{t+1} = B_1 x_t + B_0 x_{t+1} + \Theta_1 \nu_t + \Theta_0 \nu_{t+1} = B_1 x_t + B_0 x_{t+1} + u_t^1, \tag{17}$$

where $u_t^1 \equiv \Theta_1 \nu_t + \Theta_0 \nu_{t+1} = u_{t+1}^0 + \Theta_1 \nu_t$, and we assume that W_{t-1}, spanning $\sum_{l=1}^{\infty} (\Theta_{l+1} - \Theta_{l-1}) \nu_{t-l} + (B_{l+1} - B_{l-1}) x_{t-l}$, has already been partialed out. Note that u_t^1 is mechanically serially correlated (it involves both ν_t and ν_{t+1}) and correlated with u_t^0, u_{t+1}^0, and u_{t-1}^1. Moreover, without adjustment or inclusion of x_{t+1} as a regressor, x_{t+1} will appear as part of the residual and will induce further correlation between period t residuals for $h = 1$ and period $t+1$ residuals for $h = 0$.

We adopt a GLS-style transformation to eliminate serial correlation. Define

$$\tilde{Y}_t^1 = Y_t^1 - B_0 x_{t+1} - u_{t+1}^0. \tag{18}$$

This leaves

$$\tilde{Y}_t^1 = B_1 x_t + \Theta_1 \nu_t. \tag{19}$$

We can repeat this process recursively for further horizons, defining

$$\tilde{Y}_t^h = Y_t^h - B_0 x_{t+h} - \ldots - B_{h-1} x_{t+1} - u_{t+h}^0 - \ldots - u_{t+1}^{h-1} \tag{20}$$

to obtain

$$\tilde{Y}_t^h = B_h x_t + e_t^h, \;\; h = 0, 1, \ldots, H-1, \tag{21}$$

where $e_t^h \equiv \Theta_h \nu_t$ is uncorrelated with the errors in all other time periods by construction, but is correlated with time t observations at other horizons. Then, since ν_t is Gaussian, $\{e_t^h\}_{h=0}^{H-1}, \forall\, t = 1, \ldots, T$ is an $H-$dimensional Gaussian vector of correlated random variables, but is serially uncorrelated and uncorrelated across firms, as in Equation (1), once the panel dimension is reintroduced. Then it is straightforward to apply GMLR to the system of equations defined by Equation (21), for $h = 0, \ldots, H-1$, with residuals $\{e_t^h\}_{h=0}^{H-1}$. In our setting, we let $Y_t^h = \Delta\widetilde{\log k_{i,t+h}}$ and $x_t = \widetilde{\epsilon_t^m}$ be the monetary policy surprise.

The transformation operates recursively, with estimates of the impulse response coefficients and residuals at shorter horizons used to construct $\tilde{Y}_t^h$ at longer horizons. However, despite being a recursive procedure, the resulting estimates remain the true maximisers of the likelihood because the model parameters are just-identified, preserving the structure of the EM algorithm. The correction that we propose is related to Herbst and Johannsen (2024), who highlight the potential for finite-sample bias in local projections, arising from the serial correlation between included regressors in time and IRF or control coefficients at horizons less than h. While the structure of our transformation, based on $B_0 x_{t+h}$, $B_1 x_{t+h-1}, \ldots$ etc. is similar to their bias correction, our motivation is not finite sample bias, but rather efficient estimation of the Gaussian error term, to avoid misspecification of the likelihood on which identification relies in GMLR. Additionally, we assume that the monetary policy surprise is serially uncorrelated (which we verify in the data), in which case their bias expressions are identically equal to zero in our setting.

C Additional Results

C.1 Baseline: Impact RIMP with time-varying group membership

Figure C.1: RIMP heterogeneity by industry

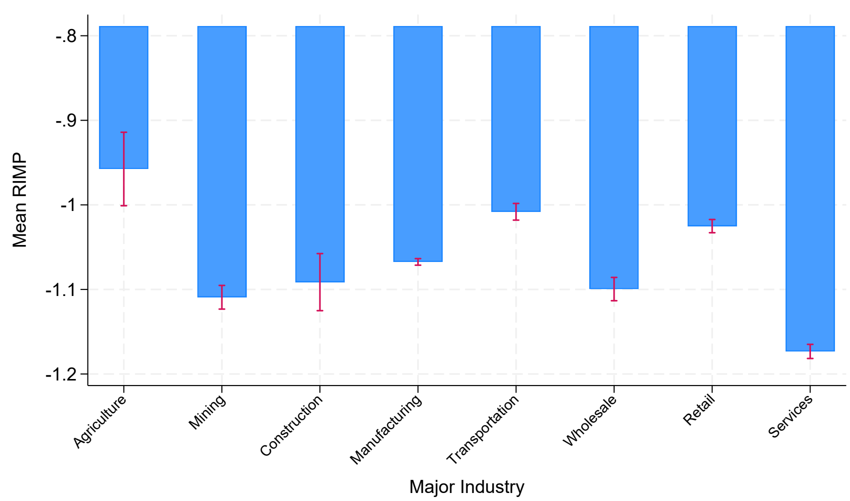

Notes: The bars plot the average predicted RIMP by major industry categories for our baseline impact RIMP estimates with time-varying group membership. 95% confidence intervals are reported in red.

Figure C.2: RIMP heterogeneity by firm age and distance to default

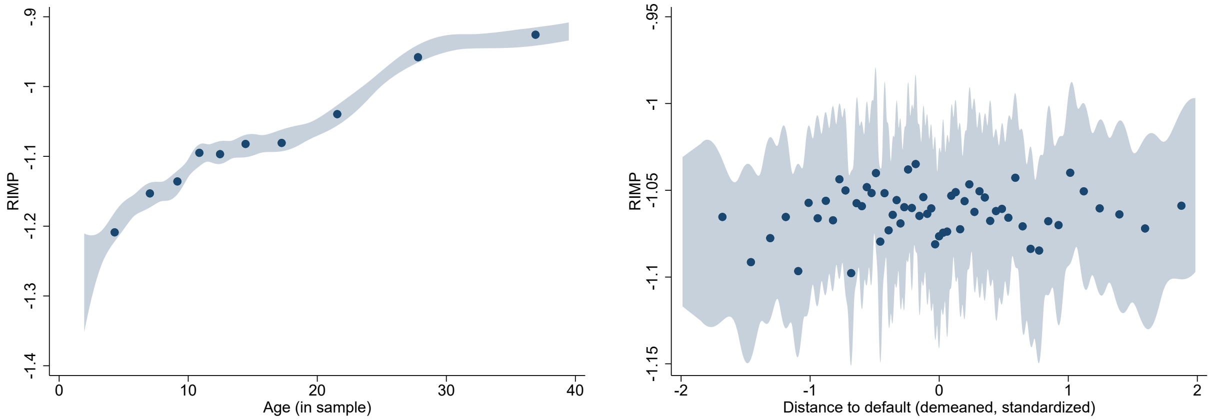

(a) RIMP and age

(b) RIMP and distance to default

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

Figure C.3: RIMP heterogeneity: LASSO

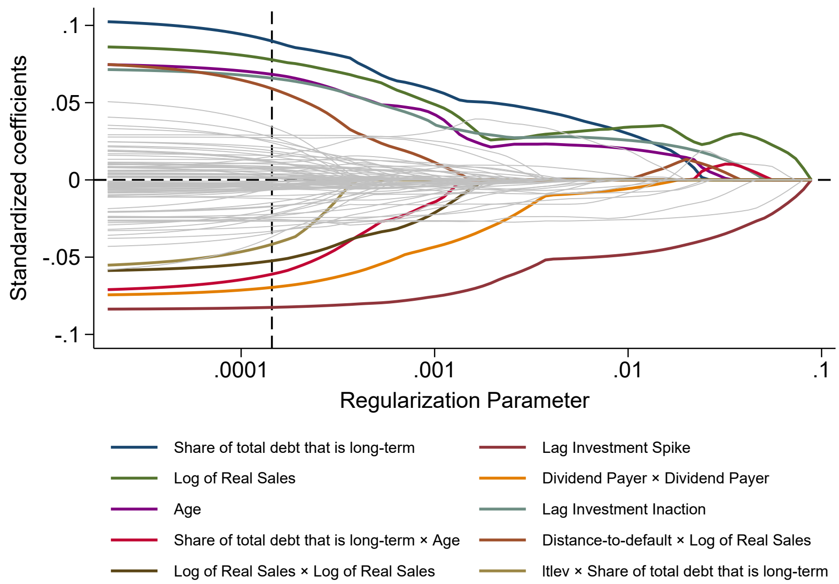

Notes: For a LASSO regression of our baseline impact predicted RIMP estimates with time-varying group membership on several predictors and their pairwise interactions at different regularization parameters, the figure plots the evolution of standardized coefficients as the regularization parameter increases. The dashed vertical line determines the optimally chosen regularization parameter; in the legend, we list the 10 largest predictors at this regularization level. Inaction is defined as having 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. "ltlev" is the ratio of total long-term debt to total assets, lagged one period. All other variables are defined in the text.

Table C.1: RIMP and observable characteristics - group-specific predictors

	Group 1 $\lambda_1 = -4.22$	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, estimated with baseline time-varying group membership. 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.

C.2 Impact RIMP with time-invariant group membership

In this appendix, we re-estimate the impact RIMP imposing time-invariant group membership. For both this specification and the estimation of local projections, the BIC falls more gradually and prominently than in our baseline estimation with time-varying group membership. In order to choose the number of groups, we therefore proceed as follows. At the firm-level, we draw 250 bootstrap samples with replacement. For each sample, we estimate our model and obtain a BIC. We then choose G as the first number of groups in which the BIC estimated in the data does not fall below the 16th percentile of BIC across bootstrapped samples for the prior group (that is, does not fall below the 68% confidence band). This approach is quite conservative and delivers $G = 4$, which we use for the results in Figure C.4 and C.5 and Table C.2. It could be argued, however, that time-invariant group membership should require a larger number of groups, given its more restrictive nature. In Figure C.6 we show that with $G = 8$ (where the BIC in $G + 1$ is only 0.04% lower than in G), the distribution fans out and more negative RIMPs are estimated. Nevertheless, our conclusions are unaffected: (i) the shape of the RIMP distribution is similar, (ii) the correlation with the baseline RIMP with time-varying group membership remains high (0.78), and (iii) the RIMP drivers are quantitatively and qualitatively similar to those shown in Table C.2.

Figure C.4: Time-invariant RIMP: predicted distribution

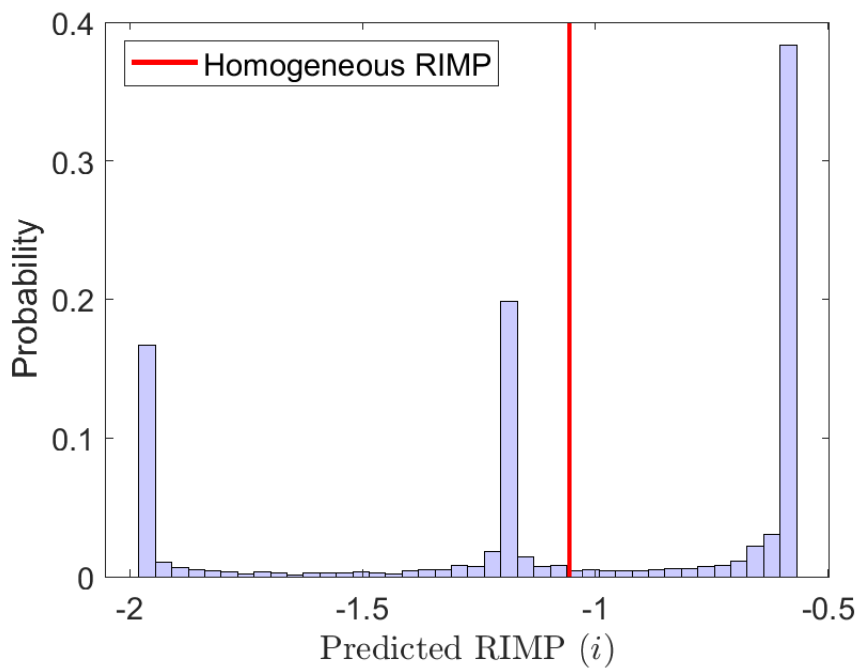

Notes: Estimates from a specification for the impact RIMP that restricts group membership to be time-invariant. The histogram plots the predicted RIMP at the firm level with $G = 4$.

Figure C.5: Time-invariant and time-varying RIMP

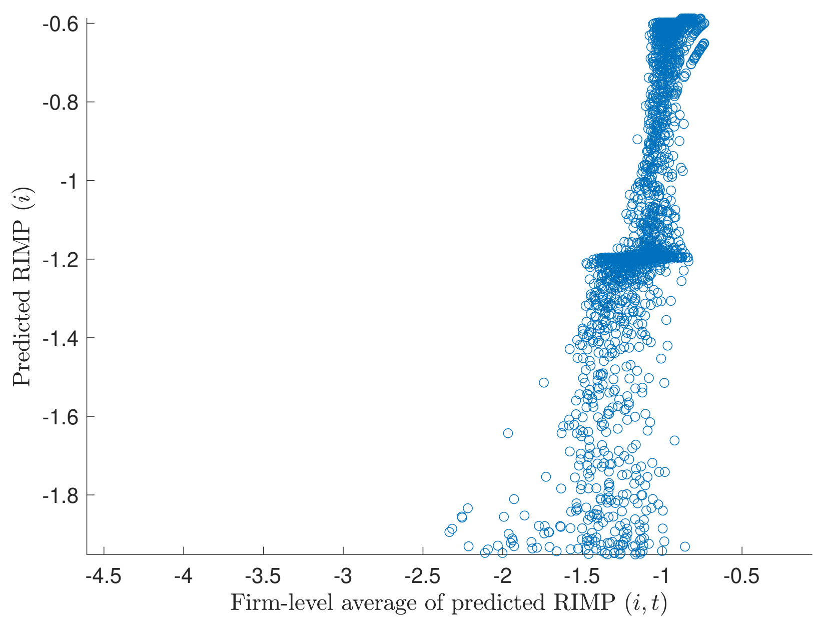

Notes: The vertical axis, measures the predicted RIMP at the firm level, estimated with time-invariant group membership, as discussed in Section 3.4. The horizontal axis averages the baseline RIMP estimates with time-varying group membership (Section 3.1), over time at the firm level. In both cases the RIMP is upon impact ($h = 0$). The correlation between the two objects is 0.81.

Table C.2: Time-invariant RIMP and observable characteristics - multivariate analysis

	RIMP(i)	RIMP(i,t)
Leverage (demeaned, standardized)	-0.006	-0.001
	(0.004)	(0.003)
Distance to default (demeaned, standardized)	0.006***	0.004**
	(0.002)	(0.002)
Share of total debt that is long-term	0.126***	0.096***
	(0.007)	(0.006)
Cash ratio	-0.429***	-0.331***
	(0.017)	(0.014)
Firm age	0.005***	0.004***
	(0.000)	(0.000)
Log of real sales (avg past year)	0.053***	0.022***
	(0.001)	(0.001)
Adjusted R^2	0.077	0.041

Notes: In the first column, the dependent variable is the predicted upon-impact ($h = 0$) RIMP estimated with time-invariant group membership. In the second column, the dependent variable is our baseline time-varying RIMP, copying column (1) of Table 3, for comparison. Robust standard errors are reported in parentheses. *,**, and *** denote significance at the 10, 5, and 1% levels respectively.

Figure C.6: Time-invariant RIMP: predicted distribution, robustness to higher G

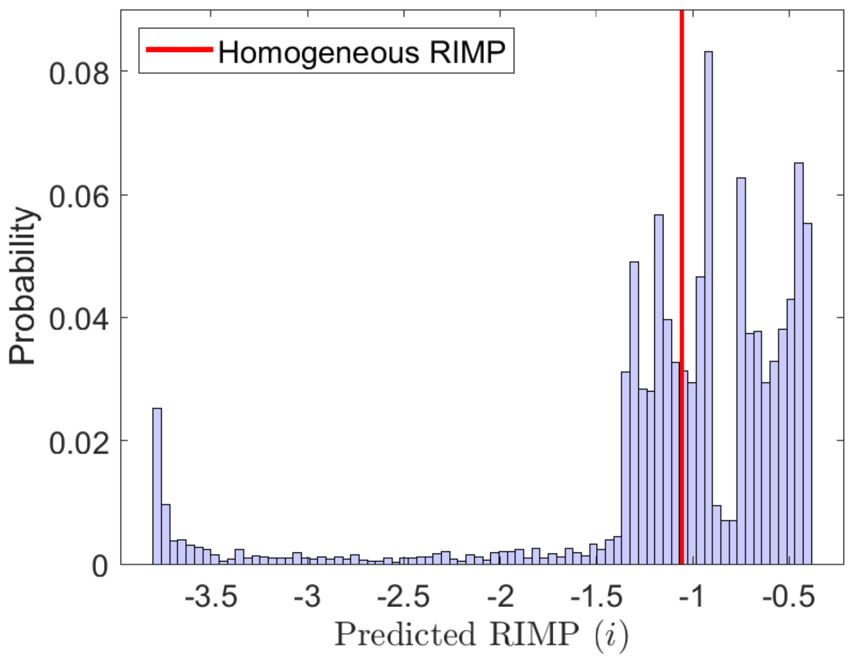

Notes: The histogram reports the distribution of the impact RIMP with time-invariant group membership at the firm level. We set $G = 8$ as discussed above to assess robustness to allowing for a greater degree of heterogeneity to compensate for the potential conservativeness of the bootstrapped BIC.

C.3 RIMP at longer horizons: additional results

Figure C.7: RIMP heterogeneity by firm size and age: longer horizons

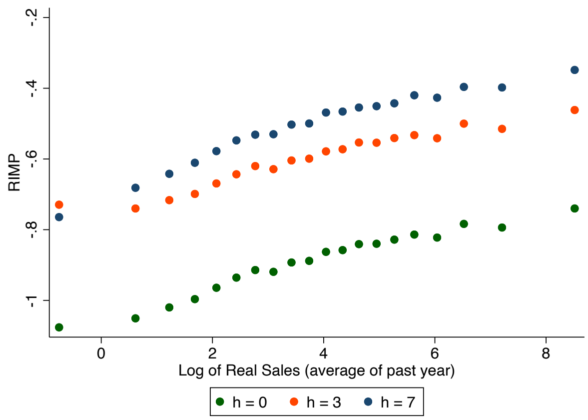

(a) RIMP and size

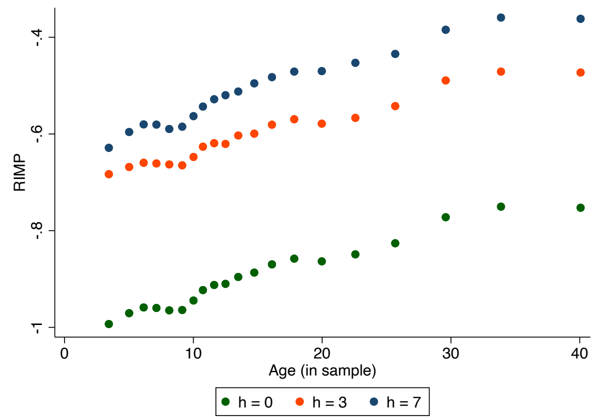

(b) RIMP and age

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). Estimates for local projection specification with $h = 0, \ldots, 7$ and $G = 4$.

Figure C.8: Predicted RIMP: impulse responses, robustness to higher G

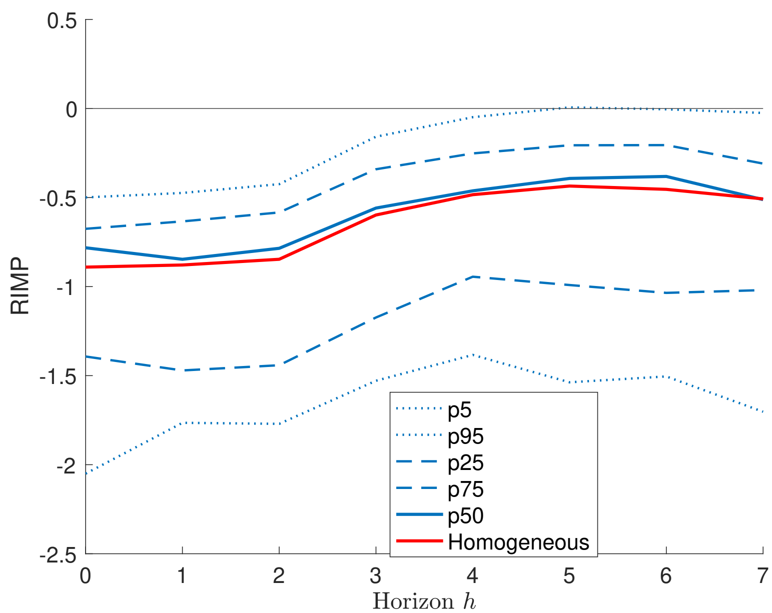

Notes: The red line ("homogeneous") plots the point estimates imposing $G = 1$. For the blue lines, we estimate the GLS-GMLR model described in Appendix C.3 with $G = 10$, compute the predicted RIMP ($\hat{\lambda}_{i,h} = \sum_{g=1}^{G} \gamma_{i,g,h} \lambda_{g,h}$) for each horizon, and then plot the quantiles (5th, 25th, 50th, 75th, and 95th percentiles) of $\hat{\lambda}_{i,h}$ across firms at each horizon. Sample size is restricted to firms with a non-missing dependent variable for all horizons.