The Investment Channel of Monetary Policy: Disentangling Firm Heterogeneity*

Thomas Drechsel[1,5] Daniel Lewis[2] Davide Melcangi[3] Laura Pilossoph[4,5]

[1]Johns Hopkins [2]UCL [3]NY Fed [4]Duke [5]NBER

NBER Summer Institue 2026

*The views expressed here do not necessarily reflect those of the Federal Reserve Bank of New York or the Federal Reserve System.

Motivation

- Widespread interest in the transmission of monetary policy to firm investment
- Particular focus on heterogeneity across firms
- Literature provides many results
  - Gertler & Gilchrist (1994): small firms more responsive to monetary policy
  - Ottonello & Winberry (2020): low-leverage firms more responsive
  - Cloyne et al (2023): young firms more responsive
  - Jeenas (2024): illiquid firms more responsive
  - ...

The conventional approach: panel local projection

Select **one** firm characteristic ex ante and include in interaction term:

$$\log k_{i,t+h} - \log k_{i,t-1} = \alpha_i^h + \boldsymbol{\beta^h x_{i,t-1} \times \epsilon_t^m} + \xi^h \epsilon_t^m + \delta^h x_{i,t-1} + \textit{controls} + \eta_{i,t}^h$$

But:

- There might be heterogeneity unrelated to the chosen $x_{i,t-1}$
- There might be heterogeneity related to many additional $x_{i,t-1}$
- Heterogeneity could be nonlinear
- Heterogeneity could be unobservable

This paper

1. Directly estimate full distribution of firms' responsiveness of investment to monetary policy (RIMP)
   - Gaussian mixture regression to assign firm-time obs to finite number of RIMPs
   - Approach successfully used for HH consumption (Lewis, Melcangi, Pilossoph, 2026)
2. Our approach
   - is *ex ante* agnostic about determinants of RIMP
   - allows a firm's RIMP to change over time
   - recovers latent and observable heterogeneity
   - can estimate *ex post* how much variation explained by a given observable
   - facilitates a horse race between competing drivers + nonlinearities

Our findings

1. Features of the RIMP distribution
   - Most responses small but non-zero
   - Small share of firm-time obs (5%) show strong responsiveness (1pp per 25bps)

2. Drivers of the RIMP distribution
   - Considerable variation within firms
   - Candidate characteristics from literature play a role
   - Also uncover new drivers
   - Large share of variation unexplained

3. Implications for structural models
   - New targets for heterogeneous firm models
   - RIMP distribution distinguishes models that one-covariate approach cannot

Empirical framework

- Follow Ottonello and Winberry (2020)
  - Same data: Compustat 1991Q2–2007Q4
  - Same left-hand-side: Δ log book value tangible capital stock
  - Same right-hand-side: HF surprises (Δ front FFF, 1-hr window)
- Homogeneous specification

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

  - Goal: heterogeneity in λ^h
  - For now: $h = 0 \rightarrow$ generalize later
  - No firm-level controls $\rightarrow$ project *ex-post*

Empirical framework (continued)

- Two preliminary steps:
  1. Absorb α_i
  2. Residualize w.r.t. aggregate controls $Z_t = [\Delta Y_t, U_t, \pi_t]'$
- Gaussian mixture linear regression model (GMLR):

$$\Delta \widetilde{\log k_{i,t+1}} = \sum_{g=1}^{G} d_{it,g} \lambda_g \widetilde{\epsilon_t^m} + \widetilde{\eta}_{i,t}$$

  - $\widetilde{\cdot}$ denotes residualized variables
  - $d_{it,g} = \mathbb{1}[it \in g]$
  - $\widetilde{\eta}_{i,t} \sim \mathcal{N}\left(0, \sigma_g^2\right)$

GMLR clustering

Complete-data likelihood is:

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

where:

- j is observation for firm i at time t
- π_g^G is unconditional probability of belonging to $g \in G$
- θ collects all parameters for $g = 1, \ldots, G$
- $\phi(\cdot; \mu, v)$ normal pdf with mean μ and variance v

d_{jg} latent: compute **expected log likelihood** over latent group membership

GMLR clustering (continued)

Expected log-likelihood is

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

where

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

is posterior probability (reflects econometrician's uncertainty)

Maximize $l(\theta)$ using **expectation-maximization (EM) algorithm**

Practical considerations

- Select $G = 4$ based on Bayesian Information Criterion
- Standard errors:
  - Fisher Information (accounts for estimation uncertainty)
  - Cluster by time (Almuzara and Sancibrian, 2024)
- Objects of interest:
  - Unconditional RIMP distribution: $\{\lambda_g, \pi_g\}_{g=1}^{G}$: consistent estimates
  - Posterior predicted RIMP distribution: $\widehat{\lambda}_j = \sum_{g=1}^{G} \gamma_{j,g} \lambda_g$
- Focus mostly on $h = 0$ today, but also construct full IRFs for each g

Results: the unconditional RIMP 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)

- Magnitude: for group 4, FFR $\uparrow$ 1 pp $\Rightarrow$ $\Delta \log k \downarrow$ by 0.5 pp

Pairwise significance

Results: the predicted RIMP distribution

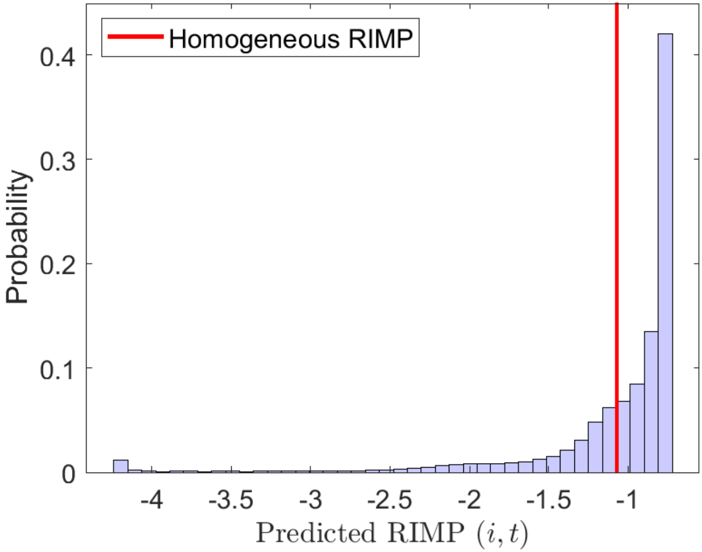

Results: persistent vs. transitory heterogeneity

- 80% of overall variance in RIMP is **within** firm (based on variance decomposition)
- Transition matrix: high RIMP status transient; low RIMP status persistent

		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

alpha

Possible drivers of the RIMP

- **"Traditional" characteristics $\rightarrow$ examine individually and jointly**
  1. Size: Small firms respond more Gertler and Gilchrist (1994)
  2. Age: Young firms respond more Cloyne, Ferreira, Froemel and Surico (2023)
  3. Default risk: Mixed results on leverage and distance to default Ottonello and Winberry (2020)
  4. Cash ratio: Liquid firms respond more Jeenas (2023)
  5. Debt maturity: Firms with more long-term debt respond less Jungherr et al. (2024)

- **"Novel" characteristics:**
  1. Perceived discount rates: larger wedge means more responsive Gormsen and Huber (2025)
  2. CFO optimism: more optimistic firms respond less Graham and Harvey (2020)

Size and age

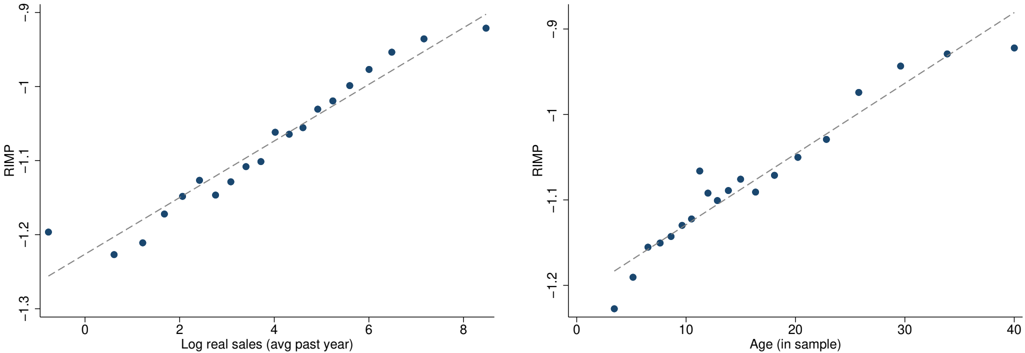

(a) RIMP and size

(b) RIMP and age

- Small firms respond **more**, as in Gertler and Gilchrist (1994)
- "Young" firms respond **more**, as in Cloyne et al (2023)

Default risk

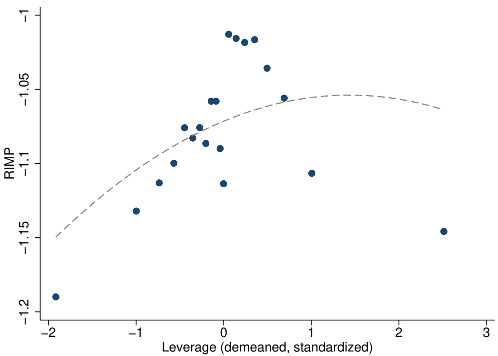

(a) RIMP and leverage

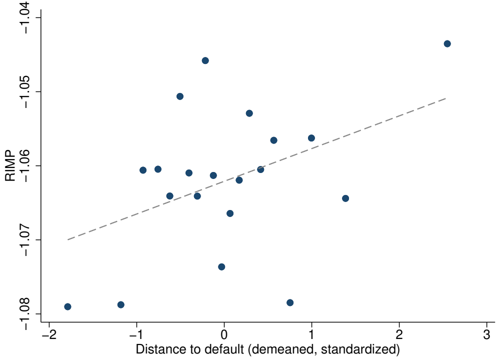

(b) RIMP and distance to default

- Low-leverage firms **more** responsive to monetary policy, as in OW (but nonlinear!)
- Distance to default ambiguous – different from OW

OW D2D

RIMP and multidimensional drivers

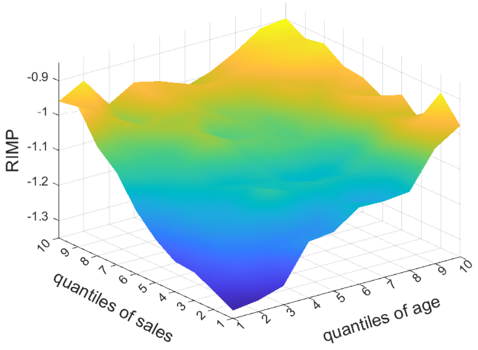

(a) RIMP by size and age

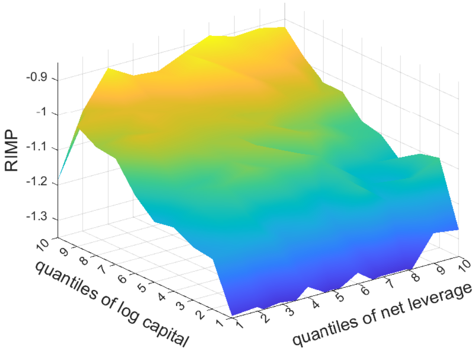

(b) RIMP by capital and net leverage ratio

- Small and young firms **more** responsive to monetary policy
- Small and less-leveraged firms **more** responsive to monetary policy

Joint 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			✓

- Majority of RIMP variation **unexplained by observables**
- Nonlinearities matter, but do not substantially increase the fit LASSO Multinomial logit

Implications of unobservable RIMP heterogeneity

- No clear benchmark for R^2
  - Unconditional investment behavior also hard to fit with firm observables

- Unobservable RIMP variation not critical for aggregate RIMP

	Implied aggregate RIMP
Using all RIMP estimates	-0.908
Using RIMP estimates predicted by leverage only	-1.062
Using RIMP estimates predicted by size only	-0.895
Using RIMP estimates jointly predicted by observables	-0.896

- Motivates the focus on newer drivers, e.g. outside of Compustat information

Discount rate wedges and stickiness

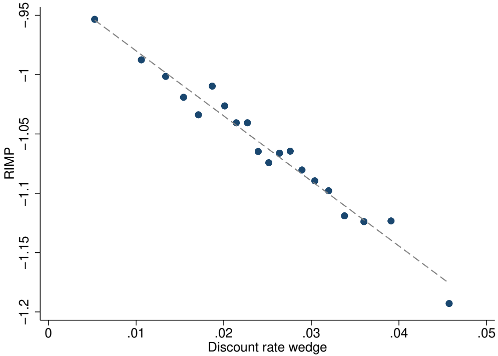

(a) RIMP and discount rate wedges

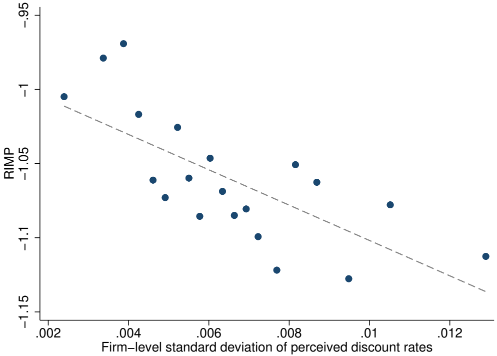

(b) RIMP and discount rate variation

- Firms with larger wedge between perceived discount rate and cost of capital (Gormsen and Huber, 2025) respond **more**
- Firms with stickier discount rates (Fukui, Gormsen, Huber, 2025) respond **less**

RIMP IRFs

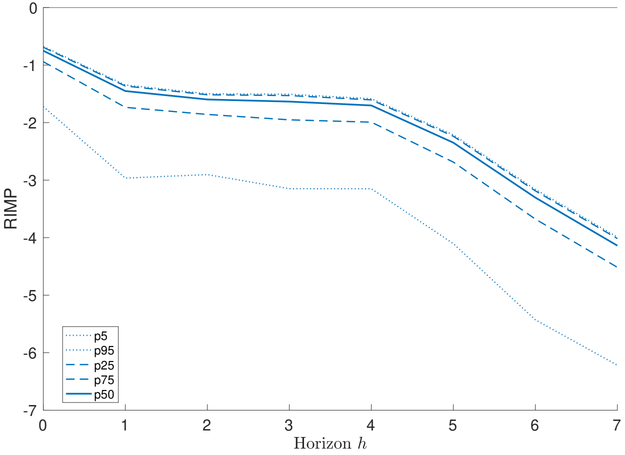

RIMP IRFs - projections on firm characteristics

	Horizon		
	0	3	7
Leverage (demeaned, standardized)	0.003	0.004	0.006
	(0.002)	(0.004)	(0.006)
Distance to default (demeaned, standardized)	0.000	0.001	0.000
	(0.001)	(0.002)	(0.003)
Share of total debt that is long-term	0.020***	0.031***	0.052***
	(0.006)	(0.009)	(0.015)
Cash ratio	−0.106***	−0.168***	−0.270***
	(0.018)	(0.029)	(0.048)
Firm age	0.001***	0.002***	0.003***
	(0.000)	(0.001)	(0.001)
Log of real sales (avg past year)	0.011***	0.017***	0.030***
	(0.004)	(0.006)	(0.009)
Time $\times$ Industry FE			
Firm FE	✓	✓	✓
Adjusted R^2	0.187	0.187	0.157

Implications for structural models

Our empirical results as a new set of targets:

1. Shape of the RIMP distribution
2. Most of RIMP variation is within firms
3. RIMP and observable drivers, multidimensional correlations

Our key argument is that (1) should be a model target:

- Show two model calibrations that match (3) but only one matches (1)

$\Rightarrow$ Full RIMP distribution is a critical target!

Model exercise

- Take **Winberry, Auclert, Rognlie, Straub (2026)** wars
  - Simplified Khan and Thomas (2013) model + NK block
  - Firm heterogeneity, capital adjustment costs, financial frictions
- Transitory, unexpected, monetary shock to the Taylor rule (SSJ perturbation)
- Compute the firm-time level GE RIMP at $h = 0$
- Two calibrations for adjustment costs (ϕ) and collateral parameters (χ):

  A. $\phi = 0.20$, $\chi = 0.20$
  B. $\phi = 0.05$, $\chi = 0.70$

RIMP moments as model targets

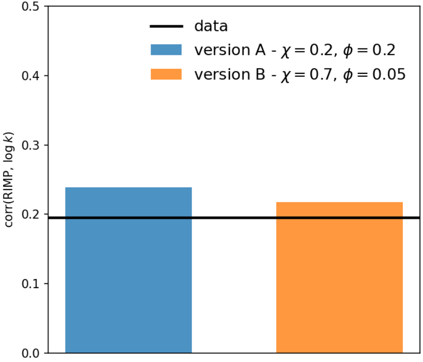

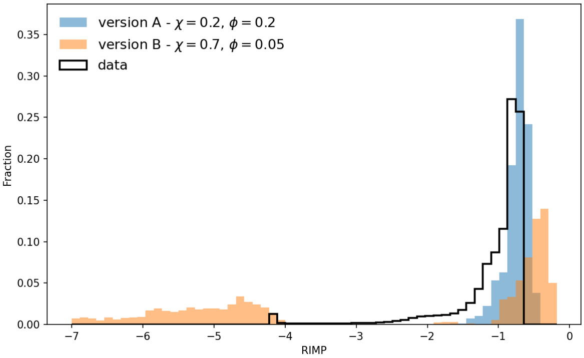

Conclusion

- We estimate distribution of firms' investment responses to monetary policy
- We uncover substantial heterogeneity:
  - Most responses are small
  - Few responses are large: up to $8x$ larger than smallest
  - Remains true at longer horizons
- We characterize RIMP heterogeneity:
  - Most variation is within firms
  - 'Old' and 'new' variables correlate with the RIMP
  - Majority of variation remains unexplained
- Full RIMP distribution is an important target for structural models

RIMP: pairwise significance

Test statistics and p-values for pairwise equality tests.

RIMP

	-4.22	-1.66	-0.90	-0.51
-4.22	4.62			
	(0.03)			
-1.66	1.48	15.53		
	(0.22)	(0.00)		
-0.90	4.62	1.48	2.88	
	(0.09)	(0.12)	(0.00)	
-0.51	4.62	1.48	2.88	3.57
	(0.06)	(0.01)	(0.03)	(0.00)

Back

Different results to OW for D2D?

Possible reasons:

- OW effect is **nonlinear**:
  - To show this, we repeat OW's specification but interact with quintiles of D2D, and find nonlinearities.
- If D2D heterogeneity depends on other forms of heterogeneity, then OW suffers from **omitted variable bias** compared to GMLR.
- Bias can also arise from higher-order relationships between x_{t-1} and errors in both investment and RIMP prediction regressions.

Back

Liquidity and debt maturity

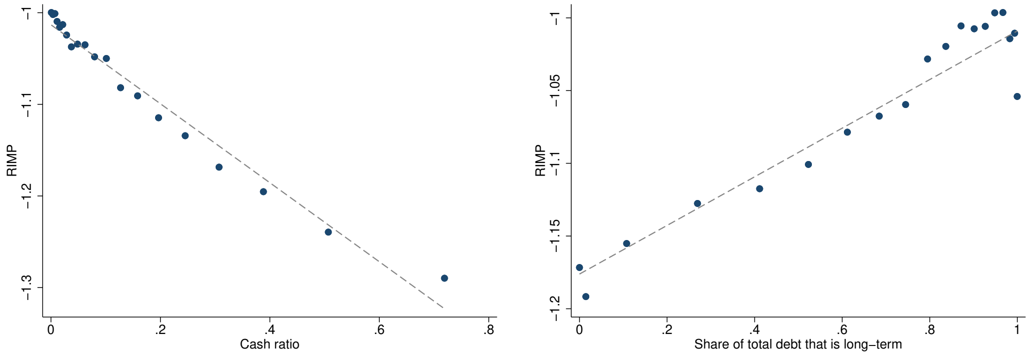

(a) RIMP and cash ratio

(b) RIMP and debt maturity

- Liquid firms respond **more** (contrasts Jeenas, 2023 for long horizons)
- Firms with more long-term debt respond **less** (Jungherr et al, 2024) Back

CFO optimism (CFO Business Outlook Survey)

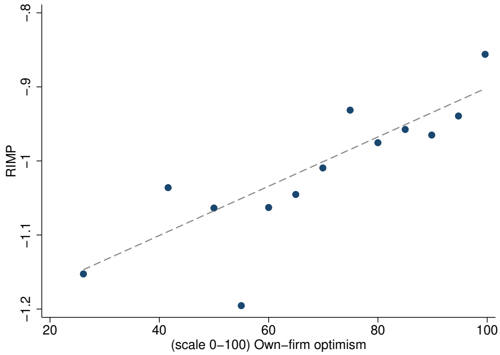

(a) RIMP and CFO optimism

- Firms with more optimistic CFOs respond **less** Back

RIMP by industry

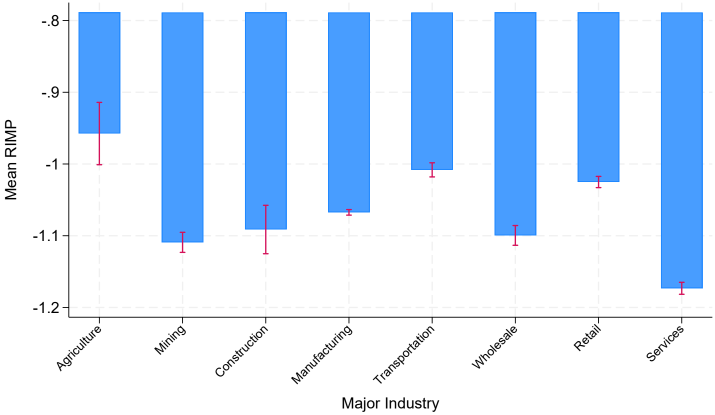

Back

LASSO regression results

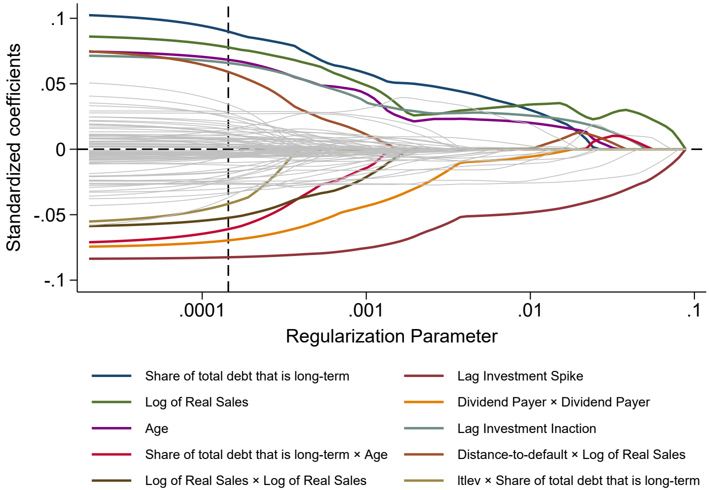

Back

Predicting individual RIMPs

	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

Back

Allowing for group-specific intercepts α_g

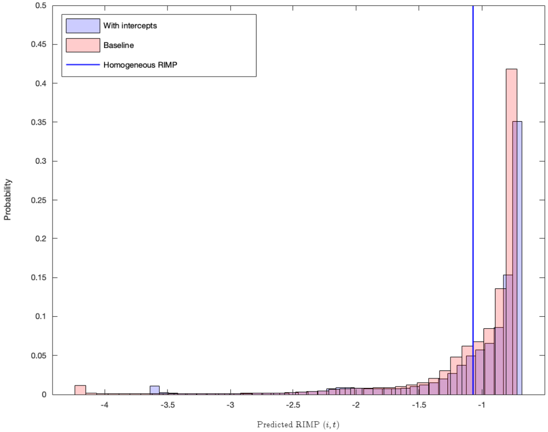

- Correlation between the predicted RIMPs is 0.99

Back

Winberry et al. (2025) wholesale producer problem

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

Other model blocks:

- Household
- Retailers (Rotemberg adjustment cost) and final-good producer
- Monetary authority (Taylor rule as in OW)

Back