Monetary Policy and Corporate Investment: Sufficient Statistics with Heterogeneous Firms*

Antoine Hubert de Fraisse† Peter Maxted‡ Kunal Sangani§
David Sraer¶

March 4, 2026

[PRELIMINARY AND INCOMPLETE—RESULTS SUBJECT TO CHANGE]

Abstract

What is the effect of monetary policy on corporate investment? To answer this question, recent contributions have exploited high-frequency monetary policy shocks aggregated at the quarterly level. We show that these shocks are endogenous to the business cycle, and thus are unlikely to be valid instruments. As an alternative, we propose a simple methodology based on Q-theory that exploits high-frequency identification but does not suffer from this bias. Empirically, our analysis uncovers a significant role for firm-level heterogeneity in driving the aggregate response of corporate investment to monetary policy.

*This is a preliminary and incomplete draft. Results and interpretations are subject to substantial revision.

†LSE; a.hubert-de-fraisse@lse.ac.uk
‡UC Berkeley; maxted@haas.berkeley.edu
§Northwestern; sangani@northwestern.edu
¶UC Berkeley; sraer@berkeley.edu

1 Introduction

How does monetary policy affect corporate investment? Despite being a central question for both macroeconomists and corporate-finance economists, the size and composition of the corporate investment response to monetary policy remain challenging to pin down empirically. The well-known difficulty is, of course, identification—the same macroeconomic forces that drive firms' investment opportunities also shape monetary policy.

The high-frequency identification approach has offered the literature a promising way forward by isolating unexpected monetary policy news using intraday interest-rate movements in narrow windows around policy announcements. However, because corporate investment is only measured at the quarterly (or annual) frequency, there is no natural way of directly studying investment in such a high-frequency framework. As an alternative, one common approach in the empirical investment literature is to aggregate high-frequency surprises at the quarterly frequency in order to estimate the effect of monetary policy on corporate investment. This paper challenges the validity of this widely-used empirical strategy. Echoing the higher-frequency predictability detailed in Bauer and Swanson [2023], we document substantial predictability in U.S. monetary policy surprises using lagged yield changes: surprise tightenings (resp. loosenings) are significantly more likely to occur during tightening (resp. loosening) cycles. This predictability complicates the causal interpretation of lower-frequency (e.g., quarterly) investment regressions, because it suggests that monetary policy surprises covary with broader macro-financial conditions that also shape investment over quarterly horizons.

In light of this problem, we propose an alternative, Q-theory-based approach to estimating the corporate investment response to monetary policy. This approach combines high-frequency identification of monetary policy's effect on firms' market values with a simple, transparent, model-based mapping from valuation changes to investment changes (i.e., Q-theory). In particular, this mapping yields a tractable decomposition of the aggregate investment response into an average-firm channel and a heterogeneity channel. Empirically, we find that a one unit surprise tightening (normalized to a one percentage point increase in the one-year Treasury yield) reduces annual corporate investment by about 0.8% of U.S. GDP, with a meaningful share of this response attributable to firm-level heterogeneity.

A large empirical literature studies how monetary policy affects a variety of outcomes by exploiting high-frequency surprises around policy announcements (i.e., market-based measures of unexpected policy news constructed from interest-rate changes in narrow windows around FOMC communications). This approach has the appealing feature that, within a tight window, it is plausible that the main new information arriving to markets is the policy communication itself, which mitigates concerns about omitted-variable bias. Motivated by this feature, recent work on corporate investment (and other real outcomes) often aggregates meeting-level surprises to the quarterly frequency in order to match the frequency of investment data. Though natural, this approach implicitly assumes that monetary policy surprises are orthogonal to other forces that drive investment at quarterly horizons. This is a stronger requirement than the local orthogonality invoked in high-frequency event studies.

Our first objective is to investigate this lower-frequency orthogonality assumption. Using the statement surprise series of Acosta et al. [2025],[1] we document substantial low-frequency predictability of U.S. monetary policy surprises, with surprise tightenings (resp. loosenings) significantly more likely to occur during a tightening (resp. loosening) cycle. Quantitatively, quarterly-aggregated statement surprises are forecastable from past yield changes with an economically large R^2: one lag of quarterly yield changes delivers an R^2 of 30%, while ten lags deliver an R^2 of almost 45%. These predictability patterns are robust to alternative measures of monetary policy surprises, and to a variety of business-cycle controls. In short, monetary policy surprises comove with, and are predictable from, variables that reflect the state of the business cycle. Consequently, estimated quarterly investment responses are difficult to interpret causally.[2]

Given these identification concerns, our second objective in this paper is to develop an alternative approach to estimating the corporate investment response to monetary policy. Our alternative approach combines two ingredients. First, although quarterly-aggregated surprises are endogenous to the business cycle, high-frequency identification remains a credible tool for estimating the immediate effects of monetary policy news in tight windows around announcements. Accordingly, we leverage high-frequency identification to estimate how firms' market values respond to monetary policy news within a narrow (two-day) window around announcements. Second, we turn to Q-theory in order to map these high-frequency changes in firms' market values into lower-frequency (annual) changes in investment. We consider a standard neoclassical investment model with a constant-returns-to-scale production function, perfect competition, and quadratic adjustment costs [Hayashi, 1982]. In this environment, a firm's investment rate is proportional to its average Q (i.e., the market value of the firm's productive assets relative to their replacement cost). Equivalently, investment is proportional to the firm's enterprise value. Since Q is a sufficient statistic for firm investment in this framework, we can combine: (i) high-frequency estimates of the response of firms' enterprise values to monetary policy surprises, with (ii) low-frequency estimates of firms' investment-Q sensitivities in order to characterize each firm's (and hence the aggregate) investment response to monetary policy.

Before continuing, we first note quickly that the Q-theory mapping from Q to investment has strong empirical support in our sample period, both in the aggregate and at the firm level.

[1] This series is available from 1994 onward as part of the U.S. Monetary Policy Event-Study Database (USMPD). A statement surprise is defined as the first principal component of the changes in money-market futures rates—collectively spanning a horizon of roughly one year—in a tight window around the release of FOMC statements.

[2] In particular, we conjecture that the business-cycle endogeneity of monetary policy surprises will bias investment estimates toward zero, since surprise tightenings are more common in periods where investment is likely rising.

Although this may seem surprising in light of a well-known literature documenting failures with Q-theory, a recent reexamination by Andrei et al. [2019] finds that the aggregate relationship between investment and Q has been remarkably tight since 1995. In our firm-level data, we similarly find strong empirical support for Q-theory: estimated investment-Q sensitivities exhibit substantial out-of-sample predictive power for firm investment, delivering R^2s of roughly 20-30%.[3]

Now detailing our Q-theory-based framework more fully, the first input that it requires is a market-based measure of firms' enterprise values (where market prices are needed to capture changes in firm value in a tight announcement window). A firm's enterprise value can be written as the (market) value of the firm's equity plus debt minus cash, where debt encompasses both bonds and bank loans.[4] To estimate high-frequency changes in enterprise values using available market-price data, we impose the approximation that the market value of bank loans does not respond to monetary policy shocks, which is reasonable given that most bank loans have floating interest rates. Denote by $W_{it} = \text{Equity}_{it} + \text{Bonds}_{it}$ the market value of a portfolio of firm i's outstanding equity and bonds. Under the approximation that the market value of bank loans does not respond to monetary policy shocks, the high-frequency change in firm i's enterprise value following a monetary policy shock dr corresponds to the high-frequency change in its combined bond-and-equity portfolio value, dW_{it}. Q-theory then yields a simple sufficient-statistics formula: letting γ_i denote firm i's investment-Q sensitivity, a monetary policy shock that changes firm i's enterprise value by dW_{it} induces a change in its investment of $dI_{it} = \gamma_i \times dW_{it}$. Aggregating across firms and scaling by aggregate portfolio value W_t then delivers the following decomposition for the aggregate investment response to a one unit monetary surprise:

$$\frac{dI_t}{W_t} \; = \; \mathbb{E}_w\left[\gamma_i\right]\mathbb{E}_w\left[\frac{dW_{it}}{W_{it}}\right] + \mathbb{C}\text{ov}_w\left(\gamma_i, \frac{dW_{it}}{W_{it}}\right),$$

where $\mathbb{E}_w$ denotes a value-weighted average (weighting by W_{it}) and $\mathbb{C}\text{ov}_w$ denotes the corresponding value-weighted covariance.

This decomposition separates two channels. The first term captures the response of the (value-weighted) average firm: monetary policy moves average enterprise value $\left(\mathbb{E}_w\left[\frac{dW_{it}}{W_{it}}\right]\right)$, which then translates into investment through the average investment-Q sensitivity $(\mathbb{E}_w\left[\gamma_i\right])$. The second term captures heterogeneity: monetary policy may disproportionately affect the valuations of firms with higher (or lower) investment-Q sensitivities. For example, if policy news hits high-γ firms more strongly, the covariance term implies an amplification of the aggregate investment response.

[3]Our empirical analysis looks only at publicly traded firms. We make no claims about Q-theory's empirical validity for smaller, private firms.

[4]More precisely, enterprise value is the market value of the firm's productive assets in our model, so we will also make adjustments for net working capital when taking the model to the data.

We then take this sufficient-statistics framework to the data. Implementing the decomposition above requires two key inputs, which are estimated in separate steps. First, we estimate the high-frequency response of firm-level enterprise values to monetary policy surprises, which provides the $\frac{dW_{it}}{W_{it}}$ inputs. Second, we estimate how firm-level investment responds to changes in valuation, which provides the γ_i inputs.

For the first set of inputs, we construct firm-level W portfolios using equity prices from CRSP and corporate bond prices from TRACE. Monetary policy surprises are taken from Acosta et al. [2025].[5] We then estimate firm-specific return sensitivities to monetary policy surprises.[6] Because firm-specific estimates are noisy, we apply standard Bayesian shrinkage toward the cross-sectional mean. In addition, we validate that the cross-sectional heterogeneity in return sensitivities we identify is not merely an artifact of sampling error by conducting an out-of-sample cross-validation exercise. Specifically, we split FOMC events into training and test sets, and find that firms that are more sensitive to monetary policy in the training sample are also more sensitive in held-out events. This provides evidence that there is persistent, economically meaningful heterogeneity in how monetary policy surprises affect firm values.

The second set of necessary inputs for the decomposition above are firm-level investment-Q sensitivities. Using annual Compustat data (1998–2023, excluding 2020), we run standard investment-Q regressions for both standard (tangible) investment and for a broader "total investment" measure that incorporates intangibles [Peters and Taylor, 2017].[7] Because these firm-specific investment-Q sensitivities are again estimated with noise, we apply Bayesian shrinkage. Moreover, we again conduct an out-of-sample cross-validation exercise to verify that firms classified as more Q-sensitive in training samples remain more Q-sensitive in held-out years.

Combining these two sets of inputs yields our central new result on heterogeneity: firms with larger investment-Q sensitivities also tend to have enterprise values that are more sensitive to monetary policy surprises. That is, the covariance term in the decomposition above is negative, $\mathbb{C}\text{ov}_w\left(\gamma_i, \frac{dW_{it}}{W_{it}}\right) < 0$. Economically, this negative covariance implies that surprise tightenings disproportionately reduce valuations for precisely the firms whose investment responds more strongly to such valuation changes. Accordingly, this firm-heterogeneity channel amplifies the aggregate investment response relative to what an average-firm-only calculation would imply.[8]

Finally, we quantify the total response of corporate investment to monetary policy. Plugging

[5]Our baseline measure is their "statement surprise" series. We also explore robustness using their "event surprise" series, which is a broader measure of FOMC policy news.

[6]In the aggregate, we find that a unit monetary policy statement surprise (normalized to correspond to a one percentage point increase in the one-year Treasury yield) generates a large and significant decline in the aggregate W portfolio of roughly 15% in a two-day window. This indicates that monetary policy news produces sharp and economically meaningful movements in firms' market valuations.

[7]Specifically, we estimate firm-level slopes γ_i from a panel regression of the investment rate I_{it}/K_{it-1} on lagged Q, including firm and industry×year fixed effects.

[8]This amplification effect is robust to alternative monetary policy surprise measures and alternative approaches to estimating investment-Q sensitivities.

our baseline estimates into the sufficient-statistics decomposition above, we find that a one unit surprise tightening reduces investment by about 0.32% of the value of the aggregate W portfolio. Since the aggregate W portfolio is roughly 2.6-times U.S. GDP in our sample, this corresponds to an annual investment reduction of around 0.8% of U.S. GDP. Roughly one-fifth of this total response is attributable to the heterogeneity (covariance) channel, with the remainder coming from the response of the value-weighted average firm.

This paper proceeds as follows. Section 2 documents the business-cycle predictability of monetary policy surprises. Section 3 lays out our Q-theory-based framework, including a sufficient-statistics decomposition into an average-firm channel and a firm-heterogeneity channel. Section 4 takes the framework to the data and quantifies the implied investment response to monetary policy. Section 5 concludes.

2 Business-Cycle Variation in Monetary Policy Surprises

This section shows that the quarterly monetary policy surprises commonly used in the literature to identify causal effects of monetary policy are not orthogonal to business-cycle conditions. Specifically, we find that surprise tightenings (resp. loosenings) are significantly more likely to happen during a tightening (resp. loosening) cycle. This finding echoes the higher-frequency results in Bauer and Swanson [2023] that monetary policy surprises are predictable using publicly available macroeconomic and financial variables observed in the quarter prior to policy announcements.[9]

Main Result. Our baseline measure of monetary policy news is the statement surprise ($stmt_t$) constructed by Acosta et al. [2025], following the approach to constructing FOMC announcement surprises of Nakamura and Steinsson [2018]. This series is available from 1994 onward and is based on high-frequency futures data from the U.S. Monetary Policy Event-Study Database (USMPD). A statement surprise is defined as the first principal component of the intraday changes in different money-market futures rates covering a horizon of roughly one year, each measured over a 30-minute window that starts 10 minutes before and ends 20 minutes after the release of an FOMC statement.[10] The resulting factor is rescaled so that a unit change corresponds to a one per-

[9]Several papers show that monetary policy surprises are predictable from information observed right before the FOMC announcement: the average federal funds rate over the calendar month before the meeting and the 12-month change in total nonfarm payroll employment using data available up to the latest release [Cieslak, 2018]; Greenbook forecasts and forecast revisions based on the most recent projections available before the announcement [Miranda-Agrippino and Ricco, 2021]; option-implied Treasury skewness averaged over the month preceding the meeting [Bauer and Chernov, 2021]; and three-month pre-meeting changes in stock prices, the yield curve slope, and commodity prices [Bauer and Swanson, 2023].

[10]Specifically, the underlying inputs consist of the implied rate surprises MP1 and MP2, constructed from federal funds futures and capturing high-frequency revisions to market expectations of the federal funds rate at the current and subsequent FOMC meetings, together with changes in short-rate expectations two to four quarters ahead,

centage point change in the one-year Treasury yield measured over the same intraday window. This factor loads similarly across each of the underlying futures rates—which collectively span a horizon of roughly one year—and therefore captures how markets revise expected short-term interest rates over the next four quarters in immediate response to an FOMC statement.[11]

Our focus in this paper is on the corporate investment response to monetary policy. Because capital expenditures are only observed at the quarterly frequency, many recent applications attempting to estimate the effect of monetary policy on corporate investment have aggregated high-frequency monetary policy surprises at the quarterly level.[12] This approach implicitly assumes that the resulting quarterly surprises are orthogonal to unobserved investment opportunities. As we show below, however, this identifying assumption is unlikely to hold.

Figure 1 presents event-study estimates of the effect of monetary policy statement surprises on the one-year Treasury yield at both daily and quarterly frequency.[13] We plot the estimated coefficients $\hat{\beta}_h$ from the following OLS regression in *event time*:

$$y^{(1)}_{t+h} - y^{(1)}_{t-1} = \alpha_h + \beta_h \cdot stmt_t + \varepsilon_{t,h}, \tag{1}$$

where $y^{(1)}_{t+h} - y^{(1)}_{t-1}$ is the change in the one-year Treasury yield measured from the end of period $t-1$ to the end of period $t+h$, and $stmt_t$ is the monetary policy statement surprise in period t. The sequence $\{\hat{\beta}_h\}^{40}_{h=-20}$ traces the event-time impulse response of the one-year Treasury yield around a monetary policy statement surprise. We set $stmt_t = 0$ on days without FOMC announcements, although the results are essentially unchanged if we restrict the sample to announcement days only. We use Huber–White robust standard errors. The sample begins in February 1994, when monetary policy statement surprises first become available, and ends in December 2023. Following standard practice in the literature, we exclude the post-9/11 FOMC announcement on September 17, 2001, as well as all announcements from March to December 2020 during the COVID-19 outbreak, leaving us with 244 monetary policy statement surprises.

Panel (a) uses daily data, and shows that a unit monetary policy statement surprise—normalized to correspond to a one percentage point high-frequency change in the one-year Treasury yield in the announcement window—leads to a 95 bps increase in the one-year Treasury yield on the day of the surprise, and an overall 160 bps increase after 20 days. The figure also exhibits pronounced pre-announcement drift: yields trend upward in the direction of the announcement from day -20 to day -7. This pattern is consistent with the finding in Bauer and Swanson [2023] that

measured using the second through fourth quarterly Eurodollar futures contracts (ED2–ED4), based on Eurodollar futures up to 2021 and SOFR futures thereafter.

[11] Statement surprises thus reflect unexpected changes in both the current policy stance and in near-term forward guidance.

[12] See Appendix A for an extensive list of references (forthcoming in future drafts).

[13] Our Treasury yield curve data come from the Federal Reserve, following the methodology of Gürkaynak et al. [2007].

recent macroeconomic and financial news, which move yields prior to FOMC meetings, predict the magnitude and sign of monetary policy surprises.

Panel (b) replicates the analysis at the *quarterly* frequency. To construct quarterly monetary policy statement surprises we sum all high-frequency statement surprises within the quarter, following standard practice in the literature. β_h in Equation (1) then captures the response of the one-year Treasury yield h quarters after the quarterly-aggregated surprise. Panel (b) reveals that quarterly statement surprises are strongly predictable using past yield changes. Appendix Figure A.1 quantifies this predictability by forecasting quarterly statement surprises with lags of quarterly yield changes: one lag forecasts $stmt_t$ with an R^2 of about 30%, while ten lags deliver an R^2 of nearly 45%.

The predictability in Panel (b) of Figure 1 extends, at a lower frequency, the finding in Bauer and Swanson [2023] that daily monetary policy surprises are predictable with publicly available macroeconomic and financial variables observed in the quarter prior to FOMC announcements. Bauer and Swanson [2023] emphasize an interpretation based on incomplete information: monetary policy responds to recent news, but fixed-income investors are uncertain about the Fed's policy rule and learn about it over time. A similar interpretation can likely explain our lower-frequency finding, as the Fed's policy stance may respond not only to short-term but also longer-term macroeconomic and financial news. Such an interpretation constitutes a rejection of full-information rational expectations, although it remains consistent with market efficiency.

Regardless of the interpretation, Figure 1 makes clear that quarterly-aggregated monetary policy statement surprises are unlikely to serve as valid instruments for identifying the effects of monetary policy on investment (or other real outcomes) at quarterly horizons. Statement surprises are significantly correlated with past yield changes, and hence with past business-cycle dynamics. As such, it is difficult to ascribe quarterly changes in investment outcomes to a monetary policy statement surprise, rather than to the continuation of pre-existing macroeconomic trends. More precisely, because surprise tightenings are significantly more likely to occur during a tightening cycle—that is, when investment is likely rising—it is natural to conjecture that such pre-trends will bias estimated monetary-policy effects toward zero.

Clarifying Scope of Endogeneity Critique. Before continuing, it is important to be precise about the scope of our critique. We *do not claim* that the business-cycle endogeneity of monetary policy surprises documented here will necessarily undermine identification in high-frequency studies of monetary policy. Indeed, the Q-theory-based methodology that we later propose will rely on estimating the high-frequency response of asset prices to monetary policy surprises. Our critique is only that once these surprises are used at lower frequencies (e.g., quarterly), they are unlikely to provide a clean source of variation for identifying the causal effects of monetary policy because they systematically comove with other business-cycle conditions.

Figure 1: Monetary Policy Statement Surprises and Interest Rates (One-Year Yield)

(a) $y^{(1)}_{t+h} - y^{(1)}_{t-1} = \hat{\beta}_h \cdot stmt_t$, Daily Frequency

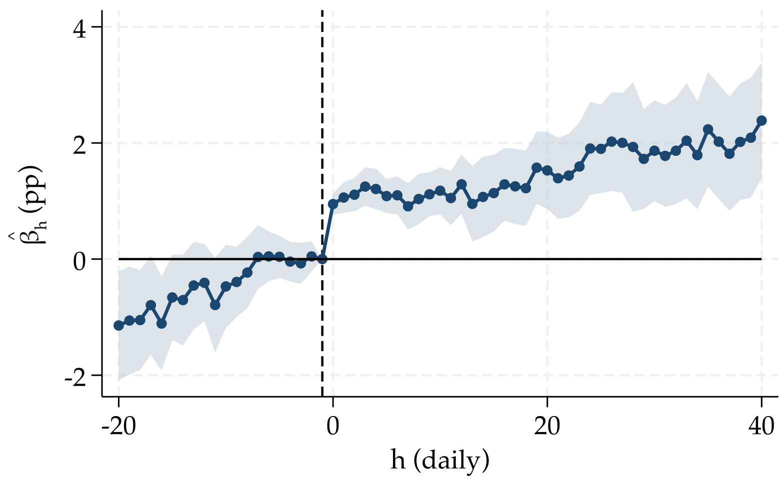

(b) $y^{(1)}_{t+h} - y^{(1)}_{t-1} = \hat{\beta}_h \cdot stmt_t$, Quarterly Frequency

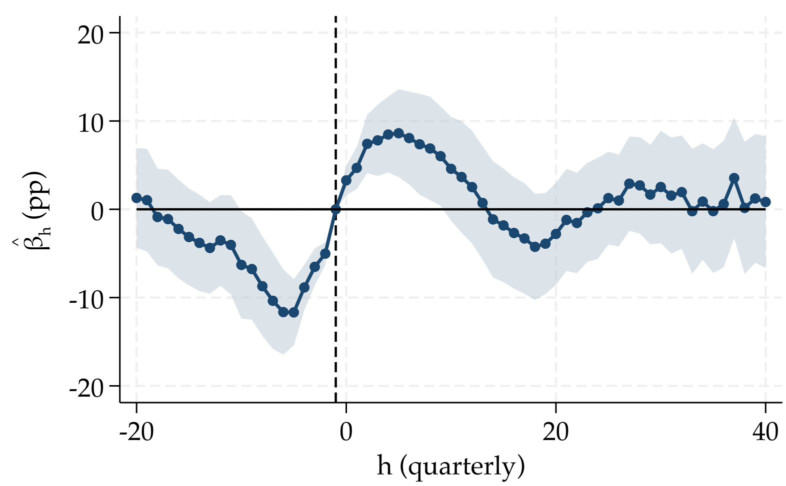

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-20}^{40}$ estimated from Equation (1). Monetary policy statement surprises come from Acosta et al. [2025]. The outcome variable is the change in the one-year $(y^{(1)}_t)$ Treasury zero-coupon bond yield. Confidence intervals are constructed at the 90% level with heteroskedasticity robust standard errors.

In more detail, a high-frequency analysis of monetary policy surprises essentially relies on the assumption that within a narrow window around the FOMC statement, the only systematic source of relevant "new" information is the statement surprise itself. By contrast, using quarterly-aggregated statement surprises to identify an investment response requires a much stronger assumption: the aggregated surprise must be orthogonal to the slow-moving macroeconomic forces that drive investment over the same horizon. Even if the surprises are only predictable ex post, the fact that surprises are strongly forecastable from lagged yield changes indicates that such surprises are systematically related to broader business-cycle dynamics.

Controlling for Observables? One simple approach to circumventing the endogeneity concerns just highlighted could, potentially, be to control for various business-cycle variables.[14] Unfortunately, Appendix Figure A.2 shows that such an approach is unlikely to adequately resolve the issue. Appendix Figure A.2 replicates the quarterly event study of Equation (1) but adds to the baseline specification, respectively, 1 lag of real GDP growth, 8 lags of real GDP growth, and 4 lags of real GDP growth, unemployment, and inflation as additional controls. Even with these controls, we continue to find that past yield changes predict the quarterly monetary policy statement surprise.

Bauer and Swanson [2023] propose an orthogonalized version of monetary policy surprises that filters out various sources of short-term predictability by regressing high-frequency monetary policy surprises on a set of pre-announcement macroeconomic and financial variables.[15] The lower-frequency predictability that we uncover is not eliminated by this procedure. Appendix Figure A.3 replicates the event study in Figure 1 but replaces the baseline statement surprise with an orthogonalized version (denoted $stmt_t^{\text{orth}}$) constructed analogously to Bauer and Swanson [2023].[16] Appendix Figure A.3 shows that this orthogonalized statement surprise remains correlated with business-cycle dynamics at lower frequencies, and is predictable using yield changes up to four quarters prior to the announcement.[17]

Overall—and as in most empirical settings—controlling for observables is unlikely to address the business-cycle endogeneity of monetary policy surprises in a satisfactory way. Once one recognizes that these "shocks" systematically comove with business-cycle conditions, adding lagged GDP growth or additional controls cannot restore orthogonality; it can only remove the compo-

[14] For instance, some papers control for lags of GDP growth, unemployment, and inflation when regressing investment on monetary policy surprises (e.g., Ottonello and Winberry [2020]).

[15] These variables include non-farm payroll surprises (i.e., deviations from consensus), employment growth, 3-month S&P 500 returns, 3-month changes in the yield curve slope, 3-month log-changes in commodity prices, and the level of skewness implied by options on 10-year Treasury note futures.

[16] Specifically, we regress the high-frequency statement surprise on the same set of pre-announcement macroeconomic and financial variables, and use the resulting residual as the orthogonalized statement surprise.

[17] An additional limitation of the orthogonalized statement surprise is that it no longer provides a *relevant* policy innovation, in the sense that it does not generate a statistically significant response of the one-year Treasury yield over the quarter in which it occurs.

nent of predictability captured by those specific controls. Because we do not know the precise sources of predictability for monetary policy surprises, it is not feasible to control for all relevant factors. Thus, residual correlation with unobserved investment opportunities will likely persist. For this reason, the remainder of the paper develops a methodology to estimate the effect of monetary policy on corporate investment that exploits high-frequency identification without relying on the exogeneity of monetary policy surprises at the quarterly horizon.

Robustness and Further Results. Our baseline approach constructs quarterly monetary policy statement surprises as the sum of all daily statement surprises within a quarter. Alternatively, Ottonello and Winberry [2020] employ a time-weighted aggregation scheme in which each daily surprise is apportioned between the current and subsequent quarter according to the fraction of the quarter remaining at the time of the announcement. Appendix Figure A.4 replicates our baseline quarterly analysis using this alternative aggregation method, and results are comparable in both cases.

We also verify that our results are robust to using the monetary policy event surprise, me_t, from Acosta et al. [2025]. These event surprises are a broader measure of FOMC policy news than statement surprises alone, as they are constructed to capture both the statement surprise and the surprise arising from the subsequent post-meeting press conference. Economically, me_t therefore incorporates additional clarification and forward-looking guidance beyond the statement itself, and tends to place somewhat greater weight on revisions to the expected policy path at longer horizons. Appendix Figure A.5 replicates our baseline quarterly event-study analysis using quarterly aggregates of me_t in place of $stmt_t$. The pattern is unchanged: quarterly event surprises remain predictable from lagged yield changes, indicating that the failure of orthogonality at the quarterly horizon also extends to broader measures of FOMC communication.

Many studies employ alternative constructions of monetary policy statement surprises. We replicate our baseline analysis using several widely used measures, including the monetary policy surprise of Nakamura and Steinsson [2018] (updated by Acosta et al. [2024]), the target and path surprise factors of Gürkaynak et al. [2005] (also updated by Acosta et al. [2024]), the monetary policy statement surprise (and its orthogonalized version) of Bauer and Swanson [2023], the monetary policy surprise of Gorodnichenko and Weber [2016], and the monetary policy surprise of Bu et al. [2021]. Appendix Figure A.6 shows that the corresponding quarterly series—aggregated in the same way as our baseline statement surprise—exhibit a similar qualitative degree of predictability from past yield changes.

Finally, although our analysis thus far has used quarterly-aggregated monetary policy statement surprises to align with the prevailing investment literature, we emphasize that the predictability we uncover is a broader low-frequency phenomenon and is not driven mechanically by time aggregation to the quarterly level. Appendix Figure A.7 provides an event study around *daily* monetary policy statement surprises using a wider window of 240 days around the surprises, and exhibits a similar pattern of predictability as our baseline quarterly event study.

3 Sufficient Statistics in a Q-Theory-Based Framework

In Section 2 we highlighted that the commonly used empirical design for estimating the corporate investment response to monetary policy—regressing investment on quarterly-aggregated monetary policy surprises—is subject to endogeneity issues. In light of these issues, we now propose an alternative, Q-theory-based methodology for estimating the corporate investment response to monetary policy.

Before beginning, it may be helpful to comment briefly on the empirical adequacy of Q-theory as a model of corporate investment. At the aggregate level, Andrei et al. [2019] document using U.S. national accounts data that the relationship between aggregate investment and Tobin's Q has become remarkably tight in the modern sample period that we will focus on empirically. For example, Andrei et al. [2019] find that an aggregate investment-Q regression has an R^2 of more than 70% over the period 1995–2015. In addition, we will show in Section 4.3 that Q-theory also has strong explanatory power for firm-level investment in our sample of U.S. publicly traded firms.[18] In particular, we find that our estimated firm-level investment-Q sensitivities predict firm investment out-of-sample with an R^2 of 20-30%.

Nonetheless, we emphasize at the outset that our Q-theory-based methodology requires a variety of strong assumptions. Thus, it is best viewed not as *the* optimal methodology for estimating the investment response to monetary policy, but rather as providing one more arrow in the quiver of applied macro- and corporate-finance economists. Given the endogeneity issues uncovered in Section 2, credible inference on monetary policy's effects will likely require the evaluation of evidence across a variety of methodologies, each with their own inherent strengths and weaknesses. One strength of our methodology is that it retains the credibility advantages of high-frequency identification (i.e., understanding how monetary policy surprises affect asset prices in a tight window around announcements), and then combines the estimated asset-price responses with a first-principles model of investment (i.e., Q-theory) in order to provide a simple, transparent, sufficient-statistics approach to identifying how investment responds to monetary policy. Our methodology also allows us to uncover novel sources of firm-level heterogeneity that will matter for aggregation, and which future research should continue to explore.

Q-Theory-Based Framework. We note quickly that the framework proposed below is, for now, more of a sketch than a fully specified model. This will be improved in future drafts (as the current draft is only preliminary).

[18] As mentioned in footnote 3, we make no claims about Q-theory's explanatory power for private firms.

Our sufficient-statistics approach is as follows. We consider a neoclassical investment model with a constant-returns-to-scale production function, perfect competition, and quadratic adjustment costs (as in Hayashi, 1982). We denote by X_{it} the firm-specific state variables of firm i at time t (e.g., capital K_{it}), by r_t the prevailing nominal interest rate at time t, and by Z_t all other aggregate state variables at time t (e.g., aggregate output). Let $V(X_{it}, r_t, Z_t)$ denote the firm's value function. The firm's first-order condition for investment is:

$$\frac{I_{it}}{K_{it}} = \delta_i + c_i^{-1}\left[\frac{\partial V(X_{it}, r_t, Z_t)}{\partial K_{it}} - 1\right] = \delta_i + \gamma_i(q_{it} - 1),$$

where $\gamma_i = \frac{1}{c_i}$, c_i is the firm's adjustment-cost parameter, δ_i is the firm's depreciation-rate parameter, and q_{it} is the firm's marginal q. With perfect competition and constant returns to scale, marginal q is equal to average Q so that:

$$\begin{aligned} \frac{I_{it}}{K_{it}} &= \delta_i + \gamma_i(Q_{it} - 1), \qquad \text{and hence} \\ I_{it} &= \delta_i K_{it} + \gamma_i\left(V(X_{it}, r_t, Z_t) - K_{it}\right). \end{aligned}$$

In the framework above, V corresponds to the firm's enterprise value (i.e., the market value of the firm's productive assets). Using the balance-sheet identity, we can relate the firm's enterprise value to the market value of its liabilities:[19]

$$\begin{aligned} V_{it} &= V(X_{it}, r_t, Z_t) = \text{Equity}_{it} + \text{Debt}_{it} - \text{Net Working Capital}_{it} \\ &= \text{Equity}_{it} + \text{Bonds}_{it} + \text{Loans}_{it} - \text{Net Working Capital}_{it}. \end{aligned}$$

To estimate how V_{it} responds to a monetary policy surprise (dr) using available market-price data, we impose the following assumption:

Assumption 1. *The market value of loans and net working capital does not respond to a monetary policy surprise:*

$$\frac{d\mathit{Loans}_{it}}{dr} \approx 0 \;\; \textit{and} \;\; \frac{d\mathit{Net\ Working\ Capital}_{it}}{dr} \approx 0.$$

The assumption that $\frac{d\text{Loans}_{it}}{dr} \approx 0$ is justified on the grounds that most bank loans have floating interest rates, as emphasized for example by Ippolito et al. [2018].[20] Of course, monetary policy

[19] In other words, the V in this framework corresponds to the market value of the firm's tangible capital and intangible assets, net of current assets.

[20] In addition, Faulkender [2005] finds in a sample of firms in the chemical industry that 89.9% of their bank loans were issued with floating rates, while Faria-e Castro and Jordan-Wood [2022] document using FR Y-14Q data that 81% of bank loans to U.S. businesses (by value) had floating interest rates at the end of 2021.

surprises may still, through general-equilibrium channels, affect the price of credit risk and hence alter the market value of bank loans, but Assumption 1 abstracts from such effects. The assumption that $\frac{d\text{Net Working Capital}_{it}}{dr} \approx 0$ is justified on the grounds that current assets and liabilities tend to be short-duration and/or floating-rate claims.[21]

Define $W_{it} = \text{Equity}_{it} + \text{Bonds}_{it}$ as the market value of firm i's portfolio of outstanding equity and bonds at time t. The benefit of Assumption 1 is that it allows us to observe the change in firm i's enterprise value V following a monetary policy surprise dr simply by looking at the change in the value of firm i's W portfolio (which is observable using market-price data):[22]

$$dV_{it} = d\text{Equity}_{it} + d\text{Bonds}_{it} + d\text{Loans}_{it} - d\text{Net Working Capital}_{it} = d\text{Equity}_{it} + d\text{Bonds}_{it} = dW_{it}.$$

Aggregating over all firms, we can then calculate the overall effect of a monetary policy shock dr on contemporaneous aggregate investment as follows (expressed as a share of the value of the aggregate W portfolio):[23]

$$\begin{aligned} \frac{dI_t}{W_t} &= \sum_i \frac{dI_{it}}{W_t} = \sum_i \gamma_i \frac{dV_{it}}{W_t} = \sum_i \gamma_i \frac{dW_{it}}{W_t} = \sum_i \left(\frac{W_{it}}{W_t}\right) \gamma_i \frac{dW_{it}}{W_{it}} \\ &= \mathbb{E}_w\left[\gamma_i\right] \mathbb{E}_w\left[\frac{dW_{it}}{W_{it}}\right] + \mathbb{C}\text{ov}_w\left(\gamma_i, \frac{dW_{it}}{W_{it}}\right), \end{aligned} \tag{2}$$

where $\mathbb{E}_w$ is a value-weighted average (using the market value of each firm's W portfolio as weights) and $\mathbb{C}\text{ov}_w$ is the value-weighted covariance.

Equation (2) breaks down the aggregate effect of monetary policy on corporate investment into two terms. The first term captures the investment response of the (value-weighted) average firm. This response is the product of: (i) the value-weighted average investment-Q sensitivity $\mathbb{E}_w\left[\gamma_i\right]$, and (ii) the average return of the value-weighted portfolio of all outstanding bonds and stocks around monetary policy shocks, $\mathbb{E}_w\left[\frac{dW_{it}}{W_{it}}\right]$. This is intuitive. If the value of the aggregate W portfolio decreases by 1% following a monetary policy tightening (i.e., $dr > 0$), then aggregate enterprise value decreases, in dollar terms, by 1% $\times W_t$. Given a weighted-average investment-Q sensitivity of $\mathbb{E}_w\left[\gamma_i\right]$, this would trigger a decrease in aggregate investment of $\mathbb{E}_w\left[\gamma_i\right] \times 1\% \times W_t$, or, relative to the value of the W portfolio, a decrease of $\mathbb{E}_w\left[\gamma_i\right] \times 1\%$.

The second term captures the role of firm-level heterogeneity.[24] As Equation (2) highlights,

[21] In our sample, the main current assets are cash and short-term investments, accounts receivable, and inventories (Days Inventory Outstanding averages about two months). The main current liability (excluding short-term debt) is accounts payable (Days Payable Outstanding also averages about two months). We again abstract from potential general-equilibrium channels by imposing Assumption 1. See Appendix D for further details.

[22] In the following equations, all derivatives are taken with respect to a unit monetary policy surprise (and hence we drop the dr to lighten notation).

[23] We clarify that aggregate investment here refers only to the investment of the publicly traded firms that we study empirically (details in Section 4), not the universe of all corporations in the U.S.

[24] The importance of covariances in micro-to-macro aggregation is well-known, and has been highlighted in recent

the importance of firm-level heterogeneity depends on the covariance between individual firms' investment-Q sensitivity, γ_i, and the sensitivity of firms' W portfolio in response to a unit monetary policy shock ($\frac{dW_{it}}{W_{it}}$). A negative covariance amplifies the aggregate response of investment to monetary policy shocks (i.e., it makes the negative investment response *more negative*), because a negative covariance implies that monetary policy shocks disproportionately affect the enterprise value of firms with high investment-Q sensitivities.

4 Taking Q-Theory-Based Framework to the Data

We now take the Q-theory-based framework from Section 3 to the data. Estimating this framework requires two (largely separate) steps. First, we need to estimate the return on firm-level W portfolios around monetary policy surprises (i.e., $\frac{dW_{it}}{W_{it}}$). Since these W portfolios are constructed using market prices, we can leverage high-frequency identification for this first step. Second, we need to estimate firm-level investment-Q sensitivities (i.e., γ_i). As discussed in Section 2, high-frequency identification is not suitable for studying corporate investment directly. Fortunately, a key advantage of imposing Q-theory is that Q becomes a sufficient statistic for investment. This allows us to recover firms' investment-Q sensitivities using straightforward, low-frequency (annual) investment-Q regressions. Then, by combining these two sets of estimates we will be able to calculate the investment response to monetary policy provided by Equation (2). We implement each of these steps in turn below.

4.1 Data Description

Return Sensitivity to Monetary Policy Surprises. When estimating the return on firm-level W portfolios around monetary policy surprises, our baseline measure of high-frequency monetary policy news continues to be the monetary policy statement surprise, $stmt_t$. For robustness, we also use the monetary policy event surprise, me_t. Both measures were detailed in Section 2, and are constructed by Acosta et al. [2025].[25]

Recall that the W portfolio corresponds to the market value of firm i's outstanding equity and bonds. For equity valuations, we use data from the CRSP US Stock Database. CRSP provides close-to-close daily equity returns, prices, and shares outstanding for U.S. publicly traded firms.

For bond valuations, we construct a firm-level panel of outstanding corporate bonds for U.S. publicly traded firms using Mergent FISD. This database contains characteristics of publicly-offered U.S. corporate bonds, including the issue date, maturity date, amount issued, principal, and coupon. It also provides a history of the amount outstanding for each bond in the sample. We

work such as Auclert [2019] and Patterson [2023].

[25] As in Section 2, we exclude all FOMC meetings between March 2020 and December 2020 to abstract from the extraordinary policy environment during the COVID-19 episode.

match this sample to data from the Trade Reporting and Compliance Engine (TRACE). TRACE provides comprehensive coverage of secondary market trades for more than 350,000 U.S. corporate bonds, including information on prices, volumes, and security identifiers. We use these data to construct daily measures of corporate bond prices at the bond level. Together with Mergent FISD, this allows us to construct a firm-level panel of outstanding corporate bond prices and returns starting in 2002. Appendix B provides details on how we merge Mergent FISD with TRACE to obtain this firm-level dataset.

Investment-Q Sensitivity. When estimating investment-Q sensitivities, we use firm-level accounting and investment data obtained from the annual Compustat database. Compustat provides detailed information on balance sheet and income statement items for the universe of publicly listed U.S. firms.

Our working sample is constructed from the Compustat annual files from 1998 to 2023 (excluding 2020).[26] We follow the same sample restrictions as in Peters and Taylor [2017]: we remove firms in regulated utilities (Standard Industrial Classification codes 4900–4999), financial firms (6000–6999), and firms categorized as public service, international affairs, or non-operating establishments (9000+). We also exclude firms with a missing or non-positive book value of assets or sales, and firms with a book value of property, plant, and equipment below $5 million.

4.2 Estimating Return Sensitivity to Monetary Policy Surprises

To bring the framework from Section 3 to the data, we start by leveraging high-frequency identification to estimate the return on firm-level W portfolios around monetary policy surprises.

Aggregate Evidence. Since this step of our analysis relies on monetary policy surprises providing a sharp and plausibly exogenous source of variation in firms' enterprise values at high frequency, we begin by verifying that this condition holds for the average value-weighted firm in our sample.

We focus on the period 2002–2023, reflecting the availability of corporate bond data required to measure changes in enterprise value. We conduct a standard event-study analysis on the value-weighted portfolio of outstanding bonds and equities issued by firms in our sample. For each FOMC meeting occurring at date t, we estimate the following day-level regression:

$$R_{t-11\to t+h} = \alpha_t + \mu_h + \sum_{\tau=-10,\tau\neq-1}^{10} \beta_\tau \mathbb{1}_{\{h=\tau\}} \times stmt_t + \epsilon_{t+h}, \tag{3}$$

[26] Although our bond-valuation data begins in 2002, we go back to 1998 when studying investment in order to make use of a slightly longer time series.

where $R_{t-11\to t+h}$ denotes the cumulative return from day $t-11$ to day $t+h$ on the value-weighted portfolio of bonds and equities in our sample, corresponding to the aggregate W portfolio introduced in Section 3.[27] Standard errors are clustered at the FOMC meeting level.

Figure 2 reports the estimated coefficients $\{\hat{\beta}_h\}$ together with their 95% confidence intervals. As Figure 2 shows, there is a pronounced decline in the aggregate W portfolio following a surprise monetary policy tightening. More precisely, between $t-1$ and $t+1$, a unit monetary policy statement surprise—normalized to correspond to a one percentage point high-frequency increase in the one-year Treasury yield—leads to a statistically significant decline of about 15% in the aggregate W portfolio.[28] This decline persists over the subsequent ten trading days.

Figure 2: Return of Aggregate Portfolio of Bonds and Stocks Around Monetary Policy Surprises

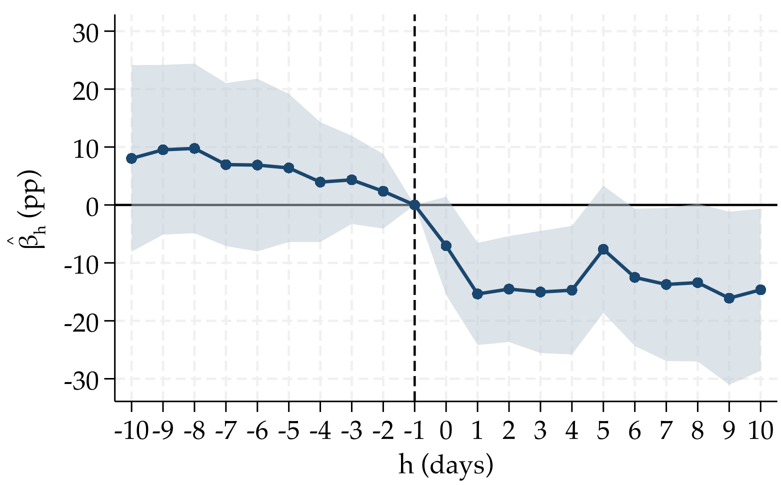

Note: Sample period 2002–2023. For each FOMC meeting taking place at date t, we estimate the following regression:

$$R_{t-11\to t+h} = \alpha_t + \mu_h + \sum_{\tau=-10,\tau\neq-1}^{10} \beta_\tau \mathbb{1}_{\{h=\tau\}} \times stmt_t + \epsilon_{t+h},$$

where $R_{t-11\to t+h}$ is the cumulative return from day $t-11$ to day $t+h$ on the value-weighted portfolio of bonds and stocks in our sample, which corresponds to the aggregate W portfolio introduced in Section 3. Monetary policy statement surprises, $stmt_t$, come from Acosta et al. [2025]. Standard errors are clustered at the FOMC meeting level.

For more, Panel (a) of Appendix Figure A.8 shows the response of the value-weighted portfolio of equities alone, estimated on the same sample and over the same period (2002–2023). The resulting pattern is nearly identical to that in Figure 2, reflecting the fact that, in our sample, equities account for roughly 91% of the market value of the aggregate W portfolio. Panel (b) of

[27] Cumulative returns are constructed from value-weighted daily returns. For equities, we use CRSP close-to-close individual stock returns. For corporate bonds, following standard practice, we compute individual bond returns from TRACE by first constructing bond-day prices as volume-weighted averages of all clean trades.

[28] Although not statistically significant, the value of the portfolio begins to decline several days prior to the FOMC meeting.

Appendix Figure A.8 extends the stock-only event study to the full sample period for which monetary policy statement surprises are available (1994–2023). The results are similar: stock values decline by about 15% between $t-1$ and $t+1$ following a surprise monetary policy tightening, and this decline persists over the ten trading days following the FOMC meeting.

Appendix Figure A.9 reproduces the same event-study analysis for the value-weighted portfolio of corporate bonds in our sample. Two differences relative to the equity response are worth noting. First, the bond portfolio exhibits a more pronounced pre-announcement drift, consistent with our findings in Section 2 that, in the days leading up to FOMC meetings, yields tend to rise prior to surprise tightenings. Second, whereas the equity response to a monetary policy statement surprise is largely realized within one day of the announcement, the bond portfolio continues to drift downward over the ten trading days following the meeting, potentially indicating a more gradual adjustment of bond prices to policy news.

We also verify that our results are robust to using the monetary policy event surprise, me_t, from Acosta et al. [2025]. Appendix Figure A.10 shows that the aggregate W portfolio exhibits a comparable decline following a surprise tightening measured by me_t. It is also worth noting that, unlike the monetary policy statement surprise, the monetary policy event surprise shows no evidence of differential pre-trends: cumulative returns prior to meetings are statistically indistinguishable from zero for equities and bonds, and there is no systematic pre-announcement drift in the direction of the subsequent tightening.

In the firm-level analysis we now turn to, we measure the response of a firm's W-portfolio value to a monetary policy surprise as the change in the value of the firm's combined bond-and-equity portfolio between the trading day preceding the FOMC meeting ($t-1$) and the trading day following the meeting ($t+1$). We adopt a two-day window as our baseline specification because the aggregate evidence in Figure 2 indicates that the cumulative price adjustment is not fully realized on the announcement day alone.[29]

Firm-Level Return Sensitivity to Monetary Policy Surprises. We now examine the return sensitivity of firm-level W portfolios to monetary policy surprises. For each firm i, we regress cumulative returns around FOMC announcements on the monetary policy surprise:

$$R_{it} = \alpha_i + \beta_i \cdot stmt_t + \epsilon_{it}, \tag{4}$$

where R_{it} is the two-day cumulative return on stocks and bonds (day 0 and day 1) around event t. We require firms to have non-missing returns for at least 50 FOMC events to estimate return sensitivities.

[29] We will also utilize a one-day window for robustness, and find that estimated return sensitivities are directionally consistent but generally smaller in magnitude.

Because individual firm estimates are noisy, we apply a standard Bayesian shrinkage toward the cross-sectional mean. Letting $\hat{\beta}_i$ denote the OLS estimate and $\hat{\sigma}_i$ its standard error, the shrunk estimate is:

$$\beta_i^{Bayes} = \lambda_i \hat{\beta}_i + (1 - \lambda_i)\bar{\beta}, \tag{5}$$

where $\lambda_i = \hat{\tau}^2/(\hat{\tau}^2 + \hat{\sigma}_i^2)$ depends on the ratio of cross-sectional variance $\hat{\tau}^2$ to estimation variance $\hat{\sigma}_i^2$. Firms with more precise estimates receive less shrinkage. Summary statistics for our estimated firm-level return sensitivities, β_i^{Bayes}, are provided in Table 1 of Section 4.4 below (first row).

Cross-Validation for Heterogeneity. To verify that we are capturing meaningful heterogeneity in return sensitivity rather than estimation noise, we validate the return sensitivity estimates using a 90/10 split of FOMC events. We sort events by the magnitude of the monetary policy shock and assign every tenth event to the test set, ensuring that the test set spans the full range of shock sizes. Using only training events, we estimate firm-level slopes β_i from Equation (4), apply Bayesian shrinkage, and sort firms into 20 bins based on their shrunk β_i^{Bayes}. We then use the held-out test events to estimate the pooled return-shock slope within each bin, controlling for firm fixed effects.

Figure 3 displays the results. The x-axis shows the average shrunk β_i^{Bayes} for firms in each bin, estimated using only training events. The y-axis shows the realized return-shock slope for the same firms, estimated using only test events that were *not used* to construct the bins. The 95% confidence intervals are based on standard errors clustered by both firm and event.[30] Figure 3 confirms that the in-sample β_i^{Bayes} estimates capture persistent heterogeneity in return sensitivity, as firms in high-β bins also exhibit significantly higher sensitivity in the held-out events.

4.3 Estimating Investment-Q Sensitivity

Baseline Firm-Level Investment-Q Sensitivity. The next step in taking the framework from Section 3 to the data is to estimate investment-Q sensitivity at the firm level. Our primary sensitivity measure comes from estimating firm-specific slopes from a panel regression of investment on Q:

$$\frac{I_{it}}{K_{it-1}} = \alpha_i + \gamma_i \cdot Q_{it-1} + \delta_{jt} + \epsilon_{it}, \tag{6}$$

where I_{it}/K_{it-1} is the investment rate, Q_{it-1} is lagged Tobin's Q, α_i are firm fixed effects, and δ_{jt} are industry×year fixed effects that absorb time-varying sector shocks (based on the Fama-French 49 industry classification).[31] We require firms to have at least 10 years of non-missing

[30]See Appendix Figure A.11 for the corresponding cross-validation using monetary policy event surprises (me).

[31]The Q-theory framework in Section 3 is in continuous time, so adjusted timing assumptions are needed here to estimate the model with annual data.

Figure 3: Cross-Validation of Return Sensitivity Estimates

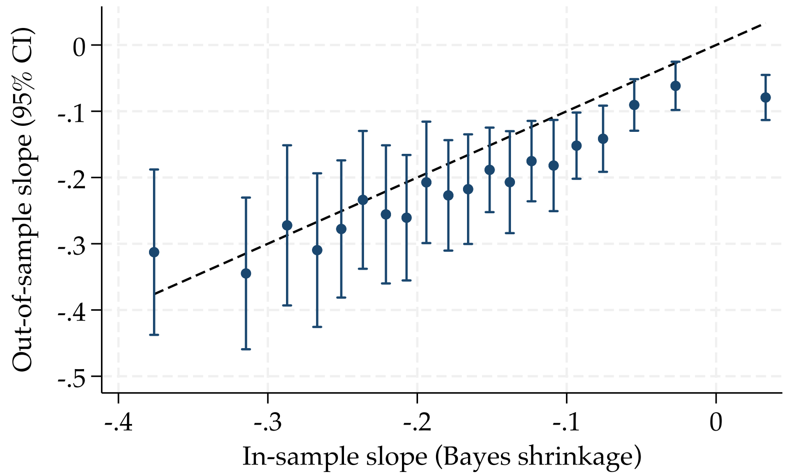

Note: This figure plots the out-of-sample return-shock slope estimated on held-out events (y-axis) against the average in-sample β_i^{Bayes} estimated on training events (x-axis) for 20 bins of firms. Error bars show 95% confidence intervals based on standard errors clustered by firm and event. The dashed line is the 45-degree line.

data to estimate investment-Q sensitivities.

We evaluate two different versions of the investment ratio in Equation (6). The first uses a tangible investment ratio (capital expenditures scaled by lagged tangible assets).[32] The second uses a total investment ratio that includes intangibles, following Peters and Taylor [2017].[33]

Because firm-specific slopes are estimated with noise, we again apply a standard Bayesian shrinkage:

$$\gamma_i^{Bayes} = \lambda_i \hat{\gamma}_i + (1 - \lambda_i)\bar{\gamma}, \tag{7}$$

where the shrinkage weight $\lambda_i = \hat{\tau}^2/(\hat{\tau}^2 + \hat{\sigma}_i^2)$ depends on the relative precision of the firm-specific estimate.

Summary statistics for our estimated firm-level investment-Q sensitivities, γ_i^{Bayes}, are again provided in Table 1 of Section 4.4 below (second and fourth rows).

Robustness: Characteristics-Based Investment-Q Sensitivity. As an alternative that avoids estimating firm-specific coefficients, we assign firms to groups based on their Fama-French 49 in-

[32]Specifically, I_{it} in this case corresponds to capital expenditures (`capx` in Compustat); K_{it-1} is the lagged physical capital stock, measured as the book value of property, plant, and equipment (Compustat item `ppegt`); Q_{it-1} is defined as the ratio of the firm's enterprise value to its physical capital stock, where enterprise value is measured as the market value of outstanding equity (Compustat items `prcc_f`$\times$`csho`), plus the book value of debt (Compustat items `dltt` + `dlc`), minus the firm's current assets (Compustat item `act`).

[33]In this case, total investment corresponds to the sum of capital expenditures plus investment in intangible capital, defined as R&D + (0.3 × SG&A); total capital corresponds to the book value of tangible assets plus the replacement cost of intangible capital provided by Peters and Taylor [2017]; total Q corresponds to the standard Tobin's Q but normalized by total capital instead of tangible capital.

dustry classification and size decile, and then estimate a single γ for each industry$\times$size cell:

$$\frac{I_{it}}{K_{it-1}} = \alpha_{g(i)} + \gamma_{g(i)} \cdot Q_{it-1} + \delta_{jt} + \epsilon_{it}, \tag{8}$$

where $g(i)$ denotes firm i's industry-size group. This approach trades off within-group heterogeneity for more precise group-level estimates, which we again shrink toward the sample mean using a similar Bayesian shrinkage. Table 1 (third and fifth rows) provides summary statistics.

Cross-Validation for Heterogeneity. To again verify that we are capturing meaningful heterogeneity in investment-Q sensitivity rather than estimation noise, we perform leave-one-year-out cross-validation. For each year t in our sample, we estimate firm-specific slopes γ_i from Equation (6) after excluding all year-t investment observations from the estimation. We apply Bayesian shrinkage, then sort firms into 20 bins based on their shrunk γ_i^{Bayes}. Within each bin, we estimate the pooled investment-Q slope using only year-t investment observations, where both investment and Q have been residualized against firm and year fixed effects. We repeat this procedure for each year and average the out-of-sample slopes across years within each bin.

Figure 4 displays the results.[34] The x-axis shows the average shrunk γ_i^{Bayes} for firms in each bin, estimated using only training data (i.e., years other than the held-out year). The y-axis shows the realized investment-Q slope for the same firms, estimated using only held-out year data that was *not used* to construct the bins. The 95% confidence intervals are constructed from robust standard errors.[35] Panel (a) of Figure 4 plots standard (tangible) investment, while Panel (b) plots total investment. In both panels, the out-of-sample slopes lie close to the 45-degree line, meaning that firms classified as having higher investment-Q sensitivity based on training data also exhibit higher investment-Q sensitivity in the held-out year. This out-of-sample test provides reassurance that our γ_i^{Bayes} estimates capture economically meaningful variation in investment-Q sensitivity across firms.

Evaluating Firm-Level Q-Theory Performance. Finally, because our sufficient-statistics framework relies on Q-theory as a model of firm investment, we now assess the extent to which our estimated γ_i^{Bayes} coefficients have out-of-sample predictive power for firm-level investment. In doing so, we provide a test of Q-theory's explanatory power for investment at the firm level.

In order to evaluate Q-theory's predictive power for firm-level investment, we perform a different version of leave-one-year-out cross-validation, as follows. For each year t in our sample, we estimate firm-specific slopes γ_i from Equation (6) after excluding all year-t investment obser-

[34] See Appendix Figure A.12 for the corresponding cross-validation using characteristics-based estimates of investment-Q sensitivity.

[35] For each bin and held-out year, we estimate the standard error of the within-bin regression slope, then average these standard errors across years and apply the usual ± 1.96 critical value.

Figure 4: Cross-Validation of Investment-Q Sensitivity Estimates

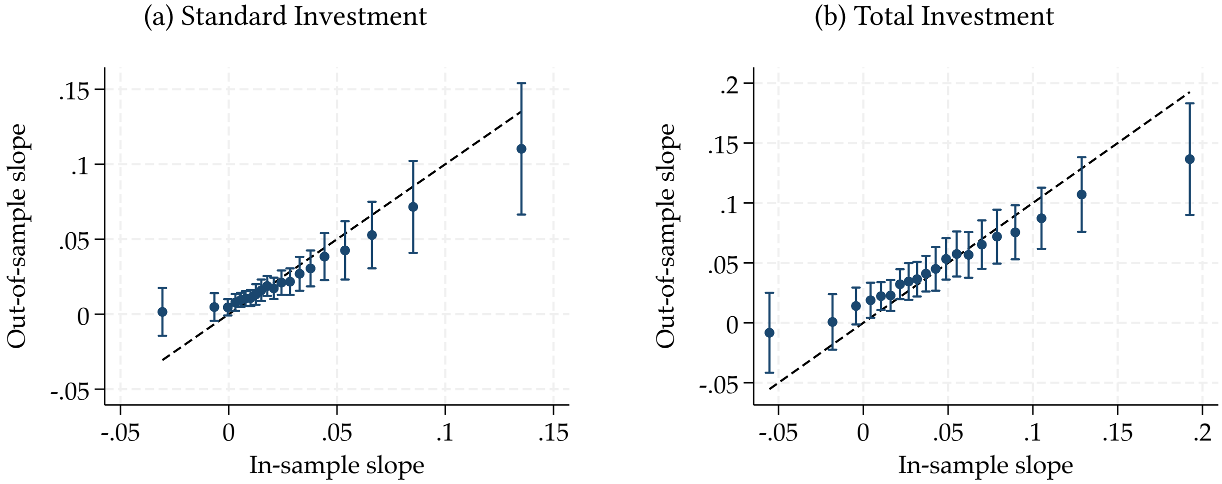

Note: Each panel plots the average out-of-sample investment-Q slope (y-axis) against the average in-sample γ_i^{Bayes} (x-axis) for 20 bins of firms. Error bars show 95% confidence intervals. The dashed line is the 45-degree line.

vations from the estimation. We require at least 10 years of non-missing data per firm and apply Bayesian shrinkage to the resulting estimates using approximate standard errors computed from the training-sample residuals. We also compute training-sample means $\bar{I}_i$ and $\bar{Q}_i$ for each firm using only the training years.

For each held-out year t, we compare the actual demeaned investment of $I_{it}/K_{it-1} - \bar{I}_i$ to the training-sample prediction for demeaned investment, $\gamma_i^{Bayes} \times (Q_{it-1} - \bar{Q}_i)$.[36] Specifically, we stack all held-out year observations and plot binned scatter plots of actual demeaned investment against predicted demeaned investment. Importantly, the γ_i^{Bayes} estimates used for predicting year-t investment are obtained from regressions that exclude year-t investment observations, ensuring a genuinely out-of-sample test.

Figure 5 displays the results. Panel (a) uses standard investment, while Panel (b) uses total investment. In both cases, the binned scatter plots reveal a tight, approximately linear relationship between predicted and actual demeaned investment with slopes that are close to one. Moreover, the out-of-sample R^2 values of 20-30% indicate substantial explanatory power. This justifies our usage of Q-theory as a simple—but nonetheless realistic—model of firm investment.

4.4 Combining Return Sensitivities with Investment-Q Sensitivities

Having discussed both the estimation of firms' return sensitivity to monetary policy surprises and firms' investment-Q sensitivity, we now combine these estimates in order to calculate the

[36]When predicting investment, we do not want to include industry×year fixed effects because doing so would require using outcomes from the held-out year.

Figure 5: Out-of-Sample Investment Predictability

(a) Standard Investment (b) Total Investment

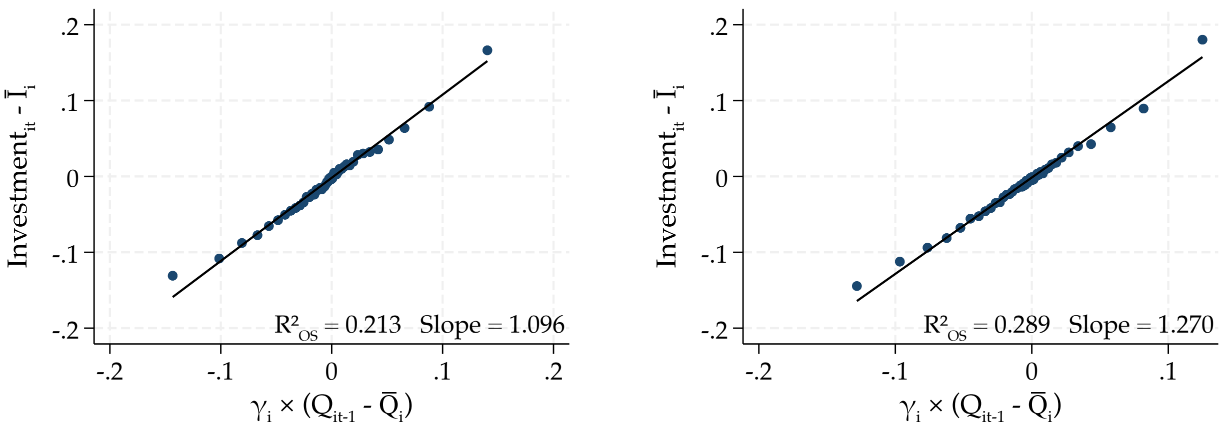

Note: Each panel plots binned averages of actual demeaned investment (y-axis) against predicted demeaned investment (x-axis). Prediction uses training-sample γ_i^{Bayes} estimates, and firms' investment and Q are demeaned using training-sample means. The fitted line is from the underlying (unbinned) regression. We report the out-of-sample R^2.

corporate investment response to monetary policy provided by Equation (2). We first present summary statistics for our firm-level sensitivity estimates, including the Equation (2) inputs of $\mathbb{E}_w[\gamma_i]$ and $\mathbb{E}_w\left[\frac{dW_{it}}{W_{it}}\right]$. We then turn to estimating the covariance term, $\mathbb{C}\text{ov}_w\left(\gamma_i, \frac{dW_{it}}{W_{it}}\right)$, which captures the role of firm-level heterogeneity. Finally, we quantify the full corporate investment response implied by Equation (2).

Our combined analysis sample consists of the intersection of firms with valid return sensitivity estimates (β) and investment-Q sensitivity estimates (γ).[37] After merging the two sets of estimates, our sample includes approximately 2,700 firms.

Summary Statistics for Firm-Level Sensitivity Estimates. Table 1 presents descriptive statistics for all variables used in the combined analysis (the estimation of which was detailed in Sections 4.2 and 4.3 above). Table 1 reports value-weighted means and standard deviations, where weights are each firm's average share of the aggregate W portfolio across FOMC events. For Equation (2), our baseline measure of the average return sensitivity $\mathbb{E}_w\left[\frac{dW_{it}}{W_{it}}\right]$ is -0.16, and our baseline measure of the average investment-Q sensitivity $\mathbb{E}_w[\gamma_i]$ is 0.02.

[37] Recall that we require firms to have non-missing returns for at least 50 FOMC events to estimate return sensitivities, and we require firms to have at least 10 years of investment and Q data to estimate investment-Q sensitivities.

Table 1: Descriptive Statistics

Variable	N	Mean	SD	Min	P25	Median	P75	Max
Return sensitivity ($stmt$)	2,693	-0.16	0.25	-0.55	-0.27	-0.20	-0.12	0.20
I-Q sensitivity (firm-level)	2,665	0.02	0.02	-0.09	0.01	0.02	0.04	0.22
I-Q sensitivity (FAMA×Size)	2,677	0.01	0.01	-0.01	0.01	0.02	0.02	0.05
I-Q sensitivity (total, firm-level)	2,673	0.03	0.04	-0.14	0.02	0.04	0.07	0.28
I-Q sensitivity (total, FAMA×Size)	2,677	0.03	0.01	-0.00	0.04	0.05	0.05	0.11

Note: This table reports summary statistics for estimated firm-level return sensitivity measures and investment-Q sensitivity measures. Mean and SD are value-weighted; percentiles are unweighted. Return sensitivity measures (β_i^{Bayes}) are estimated from firm-level regressions of W-portfolio returns on monetary policy surprises around FOMC announcements. Investment-Q sensitivity measures (γ_i^{Bayes}) are estimated from panel regressions of investment rates on lagged Q with firm and industry×year fixed effects. All estimates incorporate Bayesian shrinkage. "$stmt$" denotes the baseline monetary policy shock. "FAMA×Size" denotes group-level estimates by Fama-French 49 industry and size decile.

Firm-Level Heterogeneity: Sensitivity Covariance. Equation (2) highlights that firm heterogeneity matters for the investment response to monetary policy through a covariance term. In particular, the investment response will be amplified if the firms with a higher return sensitivity to monetary policy surprises are also the firms with higher investment-Q sensitivities.

We evaluate this covariance channel by estimating:

$$\beta_i^{Bayes} = a + b \cdot \gamma_i^{Bayes} + \epsilon_i, \tag{9}$$

where β_i^{Bayes} is firm i's (shrunk) return sensitivity to monetary policy and γ_i^{Bayes} is its (shrunk) investment-Q sensitivity. All regressions are weighted by the firm's average share of the aggregate W portfolio across FOMC events. Because both β_i^{Bayes} and γ_i^{Bayes} are estimated rather than observed, we compute standard errors via bootstrap.[38]

Before turning to regression estimates, Figure 6 presents binned scatter plots of return sensitivity against investment-Q sensitivity. We sort firms into 50 bins based on their γ_i^{Bayes} estimate and plot the value-weighted average β_i^{Bayes} within each bin against the value-weighted average γ_i^{Bayes}. The size of each point reflects the total value-weight of the firms in that bin. Panel (a) uses standard investment to measure investment-Q sensitivity, while Panel (b) uses total investment (i.e., including intangibles). Both panels reveal a clear negative relationship: firms in higher-γ bins also tend to have returns that are more sensitive to monetary policy.[39]

[38] We resample firms with replacement, re-estimate both β_i^{Bayes} and γ_i^{Bayes} (with shrinkage) in each bootstrap iteration, and run regression (9) on the bootstrapped sample. The standard deviation of the coefficient across bootstrap replications provides a standard error that accounts for the generated regressor problem.

[39] See Appendix Figure A.13 for the corresponding analysis using characteristics-based estimates of investment-Q sensitivity.

Figure 6: Return Sensitivity and Investment-Q Sensitivity

(a) Standard Investment (b) Total Investment

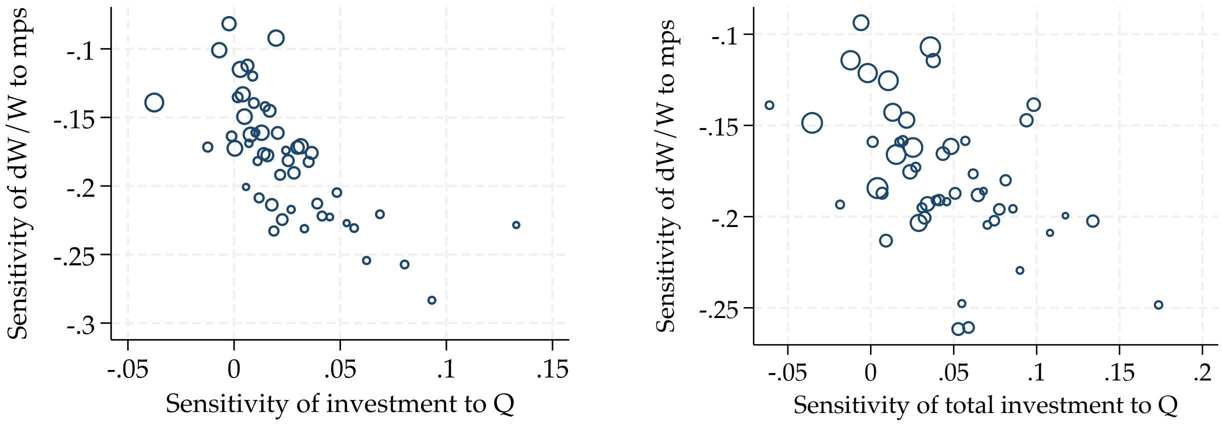

Note: Each panel plots value-weighted bin averages of return sensitivity to monetary policy (β_i^{Bayes}, y-axis) against investment-Q sensitivity (γ_i^{Bayes}, x-axis). Firms are sorted into 50 bins based on γ_i^{Bayes}. Point size reflects the total W-portfolio value of the firms in each bin.

Now turning to the regression estimates for Equation (9), Tables 2 and 3 present our main results on the role of firm-level heterogeneity in shaping the corporate investment response to monetary policy. Table 2 uses standard investment to measure investment-Q sensitivity, while Table 3 uses total investment. Column 1 of each table reports the baseline specification. The coefficient is negative and statistically significant in both tables, indicating that firms whose W portfolio is more sensitive to monetary policy surprises also have investment that is more responsive to changes in Q. The bottom row of each table reports the weighted covariance between β_i^{Bayes} and γ_i^{Bayes}, computed as $\mathbb{C}\mathrm{ov}_w(\beta_i^{Bayes}, \gamma_i^{Bayes}) = \hat{b} \times \mathrm{Var}_w(\gamma_i^{Bayes})$. Looking at Equation (2), this negative covariance implies an amplification effect stemming from firm-level heterogeneity. We quantify the implications of this finding for the transmission of monetary policy to corporate investment in Section 4.5.

The remaining columns of Tables 2 and 3 present various robustness checks. Column 2 uses one-day returns (day 0 only) instead of two-day cumulative returns to estimate return sensitivity. Column 3 replaces the firm-level γ with characteristics-based estimates, where investment-Q sensitivity is estimated at the Fama-French 49 industry$\times$size decile level. Column 4 shows our results when return sensitivity is estimated using the alternative monetary policy event surprise measure (me). The estimated coefficient remains negative and significant in all cases.

Table 2: Return Sensitivity and Investment-Q Sensitivity: Standard Investment

	(1)	(2)	(3)	(4)
	Baseline	1-day ret.	FAMA×Size	me shock
γ_i^{Bayes}	-1.11***	-0.34***	-5.79***	-0.93***
	(-12.41)	(-9.27)	(-5.81)	(-14.33)
Constant	-0.15***	-0.07***	-0.09***	-0.14***
	(-21.76)	(-26.09)	(-6.23)	(-22.31)
Observations	2,665	2,665	2,677	2,665
R^2	0.090	0.035	0.139	0.087
$100 \times \mathbb{C}\text{ov}_w\left(\beta_i^{Bayes}, \gamma_i^{Bayes}\right)$	-0.06	-0.02	-0.02	-0.05

Note: This table reports estimates from cross-sectional regressions of return sensitivity (β_i^{Bayes}) on investment-Q sensitivity (γ_i^{Bayes}) using standard investment. All regressions are value-weighted by the firm's share of the aggregate W portfolio. t-statistics in parentheses are based on bootstrap standard errors that account for the generated regressor problem. ***, **, and * denote significance at the 1%, 5%, and 10% levels. Column 1: baseline (2-day returns, firm-level γ, standard investment and standard Q, *stmt* shock). Column 2: 1-day returns. Column 3: characteristics-based γ (FAMA×Size). Column 4: β estimated with *me* shock. $\mathbb{C}\text{ov}_w(\beta_i^{Bayes}, \gamma_i^{Bayes})$ is the value-weighted covariance, computed as coefficient $\times$ $\text{Var}_w(\gamma_i^{Bayes})$.

Table 3: Return Sensitivity and Investment-Q Sensitivity: Total Investment

	(1)	(2)	(3)	(4)
	Baseline	1-day ret.	FAMA×Size	me shock
γ_i^{Bayes}	-0.45***	-0.10**	-1.87***	-0.40***
	(-5.32)	(-2.54)	(-5.20)	(-5.76)
Constant	-0.15***	-0.07***	-0.11***	-0.14***
	(-21.54)	(-25.56)	(-7.53)	(-20.96)
Observations	2,673	2,673	2,677	2,673
R^2	0.035	0.008	0.086	0.038
$100 \times \mathbb{C}\text{ov}_w\left(\beta_i^{Bayes}, \gamma_i^{Bayes}\right)$	-0.06	-0.01	-0.04	-0.05

Note: This table reports estimates from cross-sectional regressions of return sensitivity (β_i^{Bayes}) on investment-Q sensitivity (γ_i^{Bayes}) using total investment. All regressions are value-weighted by the firm's share of the aggregate W portfolio. t-statistics in parentheses are based on bootstrap standard errors that account for the generated regressor problem. ***, **, and * denote significance at the 1%, 5%, and 10% levels. Column 1: baseline (2-day returns, firm-level γ, total investment and total Q, *stmt* shock). Column 2: 1-day returns. Column 3: characteristics-based γ (FAMA×Size). Column 4: β estimated with *me* shock. $\mathbb{C}\text{ov}_w(\beta_i^{Bayes}, \gamma_i^{Bayes})$ is the value-weighted covariance, computed as coefficient $\times$ $\text{Var}_w(\gamma_i^{Bayes})$.

4.5 Quantification

Finally, we are now ready to quantify the effect of monetary policy on corporate investment implied by Equation (2). Recall that:

$$\frac{dI_t}{W_t} = \mathbb{E}_w\left[\gamma_i\right]\mathbb{E}_w\left[\frac{dW_{it}}{W_{it}}\right] + \mathbb{C}\text{ov}_w\left(\gamma_i, \frac{dW_{it}}{W_{it}}\right).$$

Using our baseline specification (2-day returns, firm-level γ, standard investment, *stmt* shock), we find from Table 1 that $\mathbb{E}_w\left[\gamma_i\right] = .0157$ and $\mathbb{E}_w\left[\frac{dW_{it}}{W_{it}}\right] = -.164$, and from Table 2 that $\mathbb{C}\text{ov}_w\left(\gamma_i, \frac{dW_{it}}{W_{it}}\right) = -.00065$, so that:

$$\frac{dI_t}{W_t} = -.164 \times .0157 - .00065 = -.32\%.$$

Our analysis therefore implies that a one unit monetary policy tightening reduces investment by about 0.32% of the value of the aggregate W portfolio. Because the market value of the W portfolio is roughly 2.6 times U.S. GDP, this decline corresponds to an annual investment reduction of about 0.8% of U.S. GDP. Of this effect, we find that about 20% is driven by the firm-level heterogeneity characterized by the covariance term.

5 Conclusion

This paper makes two main contributions. First, we illustrate that a common empirical methodology for estimating the investment response to monetary policy—regressing investment on quarterly-aggregated monetary policy surprises—is subject to endogeneity issues. Second, we develop and implement a novel Q-theory-based framework for estimating the corporate investment response to monetary policy. Our baseline estimate implies that a one unit monetary policy tightening leads to a decline in annual investment of roughly 0.8% of GDP. Future research should investigate why the firms with a higher return sensitivity to monetary policy are also the firms with larger investment-Q sensitivities, as this covariance is central to the heterogeneity-driven amplification effect that we uncover.

References

Miguel Acosta, Connor M. Brennan, and Margaret M. Jacobson. Constructing high-frequency monetary policy surprises from SOFR futures. Finance and Economics Discussion Series (FEDS) Paper 2024-034, Board of Governors of the Federal Reserve System, 2024.

Miguel Acosta, Andrea Ajello, Michael Bauer, Francesca Loria, and Silvia Miranda-Agrippino. Financial market effects of FOMC communication: Evidence from a new event-study database. Working Paper 2025-30, Federal Reserve Bank of San Francisco, 2025.

Daniel Andrei, William Mann, and Nathalie Moyen. Why did the q theory of investment start working? *Journal of Financial Economics*, 133(2):251–272, 2019.

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

Michael D. Bauer and Mikhail Chernov. Interest rate skewness and biased beliefs. Working Paper 28954, National Bureau of Economic Research, 2021.

Michael D. Bauer and Eric T. Swanson. A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37(1):87–155, 2023.

Yu Bu, Zexi Liu, and Yunzhi Xu. Interest rate news and credit supply: Firm-level evidence from the term structure of interest rates. *Journal of Monetary Economics*, 123:27–46, 2021.

Anna Cieslak. Short-rate expectations and unexpected monetary policy shocks. *The Review of Financial Studies*, 31(7):2599–2634, 2018.

Jens Dick-Nielsen. Liquidity biases in TRACE. *Journal of Financial Economics*, 94(3):439–463, 2009.

Jens Dick-Nielsen. How to clean enhanced TRACE data. SSRN Working Paper 2337908, 2014.

Chuck Fang. Monetary policy amplification through bond fund flows. SSRN Working Paper 3963516, Jacobs Levy Equity Management Center for Quantitative Financial Research, 2025.

Miguel Faria-e Castro and Samuel Jordan-Wood. How changing interest rates affect variable-rate loans to U.S. firms. Technical report, Federal Reserve Bank of St. Louis, 2022.

Michael Faulkender. Hedging or market timing? Selecting the interest rate exposure of corporate debt. *The Journal of Finance*, 60(2):931–962, 2005.

Yuriy Gorodnichenko and Michael Weber. Are sticky prices costly? Evidence from the stock market. *American Economic Review*, 106(1):165–99, 2016.

Refet S. Gürkaynak, Brian P. Sack, and Eric T. Swanson. Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *Finance and Economics Discussion Series*, 2005(66), 2005.

Refet S. Gürkaynak, Brian Sack, and Jonathan H. Wright. The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics*, 54(8):2291–2304, 2007.

Fumio Hayashi. Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, 50(1):213–224, 1982.

Filippo Ippolito, Ali K. Ozdagli, and Ander Pérez-Orive. The transmission of monetary policy through bank lending: The floating rate channel. *Journal of Monetary Economics*, 95:49–71, 2018.

Silvia Miranda-Agrippino and Giovanni Ricco. The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics*, 13(3):74–107, 2021.

Emi Nakamura and Jón Steinsson. High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics*, 133(3):1283–1330, 2018.

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

Christina Patterson. The matching multiplier and the amplification of recessions. *American Economic Review*, 113(4):982–1012, 2023.

Ryan H. Peters and Lucian A. Taylor. Intangible capital and the investment-q relation. *Journal of Financial Economics*, 123(2):251–272, 2017.

Bektemir Ysmailov. Interest rates, cash and short-term investments. *Journal of Banking & Finance*, 132:106225, 2021.

A Literature Review

A fuller set of citations is forthcoming in future drafts.

B Details on Data Construction

We combine Mergent FISD, the WRDS TRACE Enhanced Clean file, Compustat, and the Bond-Compustat linker from Fang [2025] to obtain a firm-bond-day dataset with bond-level information on outstanding amounts, coupon payments, and secondary market prices.

The sample covers all corporate and corporate-linked debt instruments in FISD observable in TRACE. This includes corporate debentures and notes (including medium-term notes and retail notes), short-term instruments (commercial paper and corporate notes), zero-coupon and payment-in-kind bonds, callable bonds, covered and other structured corporate debt, and debt-like preferred and hybrid capital securities (including preferred stock–type issues and trust preferred/capital securities). We exclude bonds issued in foreign currency.

We start from the full Mergent FISD database, using the `ISSUE`, `REDEMPTION`, `COUPON`, and `AMT_OUT_HIST` tables to recover bond characteristics and the evolution of amounts outstanding. For each bond, we construct a time series starting from the offering date and ending at maturity date. The resulting panel is merged with the `AMT_OUT_HIST` table to attach historical information on outstanding amounts. The outstanding amount series is forward-filled between reporting dates within issue; when outstanding is missing in the issuance month it is anchored to the offering amount, and negative outstanding entries in subsequent action records are recoded to zero.

The bond panel is linked to contemporaneous firm identifiers using the monthly linker developed by Fang [2025]. The linker maps issuer CUSIPs to Compustat GVKEYs in real time, explicitly accounting for subsidiary issuance, time-varying parent–subsidiary relationships, and cases in which bond issuer identifiers differ from equity issuer identifiers. See Fang [2025], for details on the construction and coverage of the linker.

The firm-bond-day panel is matched by bond CUSIP with secondary-market prices obtained from the WRDS TRACE Enhanced Clean files, which implement the transaction-level cleaning procedure described in Dick-Nielsen [2009, 2014], notably addressing cancellations, corrections, reversals, and double counting, and harmonize reporting across the February 6, 2012 TRACE system change. Cleaned transactions are aggregated to the bond–day level by computing trade-size–weighted average prices by execution date. Observations dated prior to the bond's offering date are excluded. These daily prices are merged with FISD bond identifiers and coupon information, in order to reconstruct each bond's coupon schedule and compute accrued interest on each trading day. Fixed-coupon bonds with missing or inconsistent coupon information are dropped, as accrued interest cannot be reliably computed for these observations. Accrued interest is computed per $100 of par using settlement-date accrual (T+2 with weekend adjustment), the bond's stated day-count convention when available, and explicit treatment of irregular first coupon periods, defaulting to a 30/360 convention when the day-count basis is not reported. Combining accrued interest with daily clean prices yields full bond prices at the bond-day level. Missing daily prices within a bond's life are forward-filled to measure the size of the market value of firms' outstanding bond debt and are not used to compute returns.

C Additional Figures

Figure A.1: Predictability of Quarterly Monetary Policy Statement Surprises: Cumulative R^2

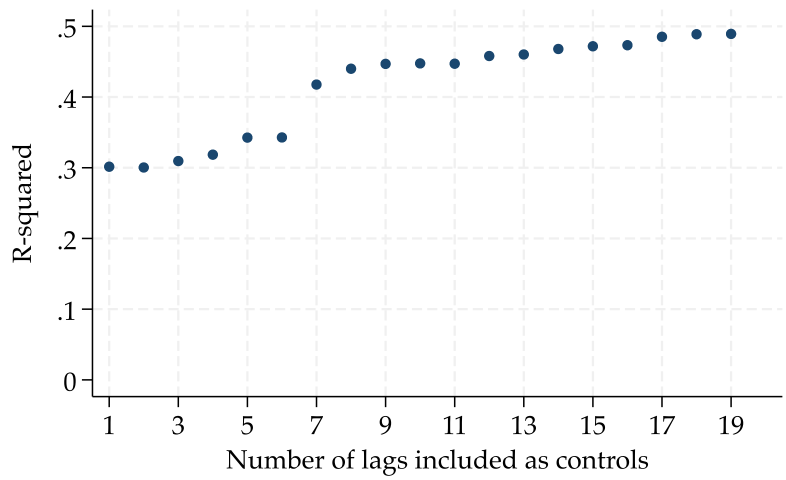

Note: We estimate the predictive regressions: $stmt_t = \alpha + \sum_{l=1}^{k} \beta_l (y_{t-l}^{(1)} - y_{t-l-1}^{(1)}) + \epsilon_t$ using various lags up to k=20. $stmt_t$ is the monetary policy surprise in quarter t, which is the sum of all the monetary policy statement surprises from Acosta et al. [2025] over quarter t. $y_t^{(1)}$ is the one-year Treasury zero-coupon bond yield in quarter t. The figure reports the R^2 of this regression against k, the number of lags used to forecast the monetary policy surprise.

Figure A.2: Monetary Policy Statement Surprises and Interest Rates: Controlling for GDP Growth

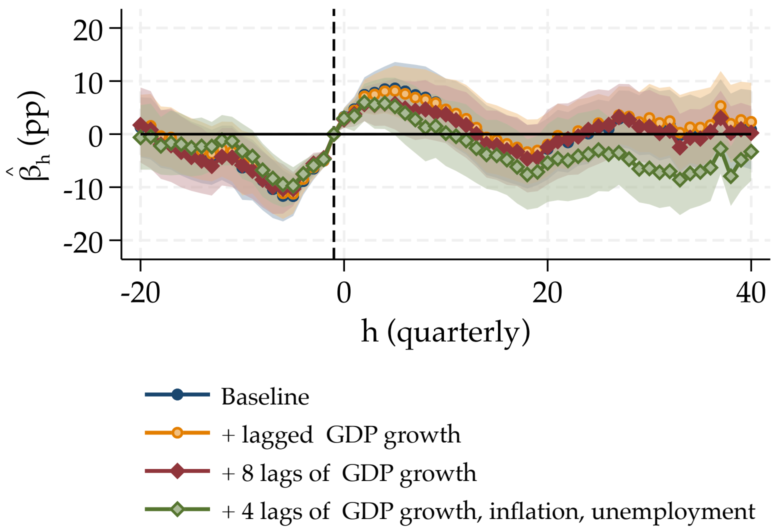

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-20}^{40}$ estimated from: $y_{t+h}^{(1)} - y_{t-1}^{(1)} = \alpha_h + \beta_h \cdot stmt_t + \sum_{l=1}^{k} \phi_h^l \text{GDP Growth}_{t-l} + \varepsilon_{t,h}$. Monetary policy surprises $stmt_t$ come from Acosta et al. [2025]. $y_t^{(1)}$ is the one-year Treasury zero-coupon bond yield in quarter t. Confidence intervals are constructed at the 90% level with heteroskedasticity robust standard errors. The blue line corresponds to the baseline case with no control for lagged GDP growth. The yellow line adds one lag of GDP growth ($k = 1$); the red line adds 8 lags of GDP growth ($k = 8$); the green line adds 4 lags of GDP growth, inflation, and unemployment.

Figure A.3: Orthogonalized Monetary Policy Statement Surprises and Interest Rates

(a) $y^{(1)}_{t+h} - y^{(1)}_{t-1} = \hat{\beta}_h \cdot stmt^{\text{orth}}_t$, Daily Frequency

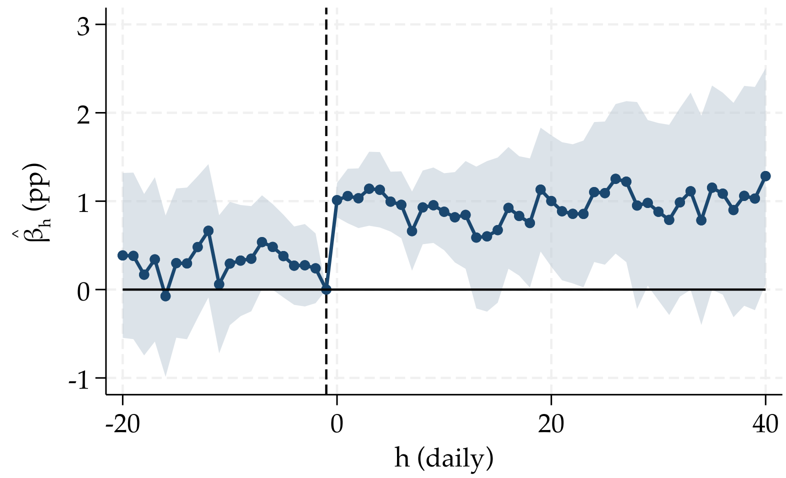

(b) $y^{(1)}_{t+h} - y^{(1)}_{t-1} = \hat{\beta}_h \cdot stmt^{\text{orth}}_t$, Quarterly Frequency

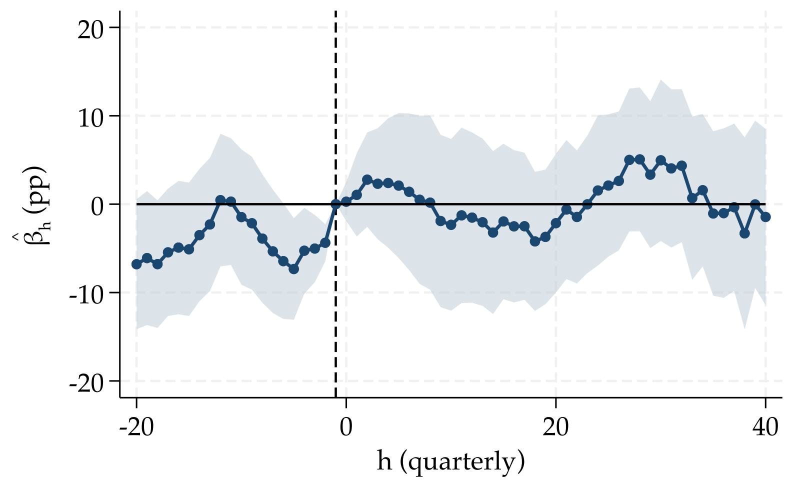

Note: This figure plots the sequence $\{\hat{\beta}_h\}^{40}_{h=-20}$ estimated from Equation (1) where we use an *orthogonalized* version of the monetary policy statement surprises from Acosta et al. [2025]. Panel (a) uses daily data and daily surprises; Panel (b) uses quarterly data and quarterly aggregated surprises. The outcome variable is the change in the one-year $(y^{(1)}_t)$ Treasury zero-coupon bond yield. Confidence intervals are constructed at the 90% level with heteroskedasticity robust standard errors.

Figure A.4: Monetary Policy Statement Surprises and Interest Rates: Time Weighting

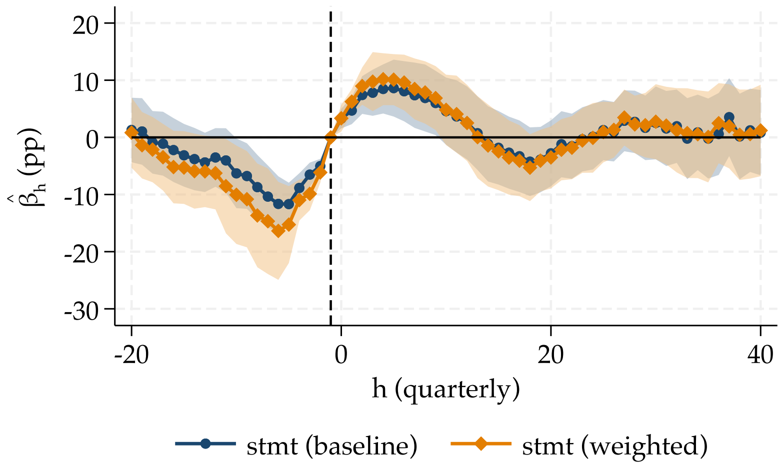

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-20}^{40}$ estimated from Equation (1). Monetary policy statement surprises come from Acosta et al. [2025]. They are aggregated within a quarter using either an unweighted sum, or a weighted-average using the fraction of the quarter remaining as weights (see Ottonello and Winberry, 2020). The outcome variable is the change in the one-year $(y_t^{(1)})$ Treasury zero-coupon bond yield. Confidence intervals are constructed at the 90% level with heteroskedasticity robust standard errors.

Figure A.5: Monetary Policy *Event* Surprises and Interest Rates

(a) $y^{(1)}_{t+h} - y^{(1)}_{t-1} = \hat{\beta}_h \cdot me_t$, Daily Frequency

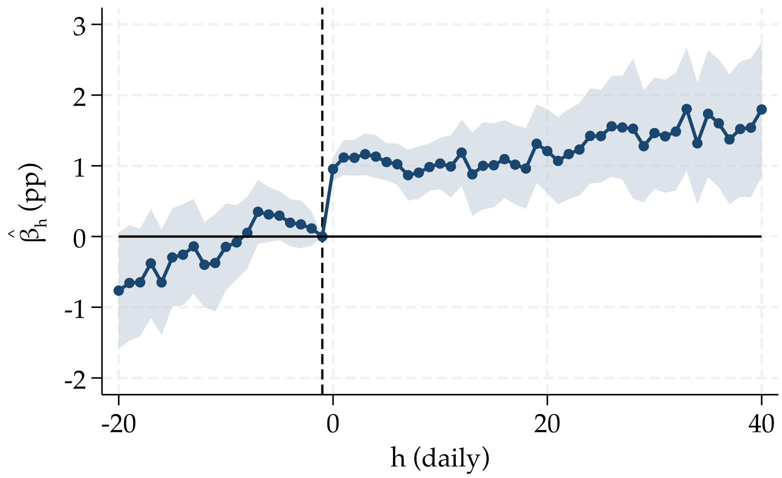

(b) $y^{(1)}_{t+h} - y^{(1)}_{t-1} = \hat{\beta}_h \cdot me_t$, Quarterly Frequency

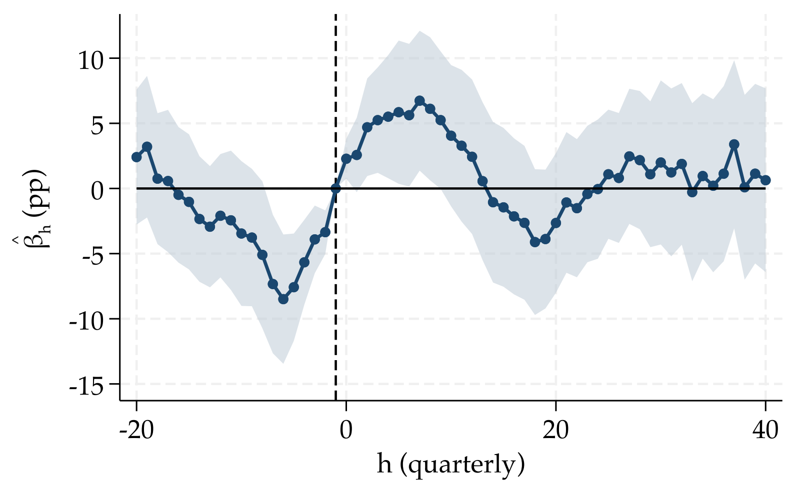

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-20}^{40}$ estimated from Equation (1) where we use the monetary policy *event* surprises from Acosta et al. [2025]. Panel (a) uses daily data and daily surprises; Panel (b) uses quarterly data and quarterly aggregated surprises. The outcome variable is the change in the one-year $(y^{(1)}_t)$ Treasury zero-coupon bond yield. Confidence intervals are constructed at the 90% level with heteroskedasticity robust standard errors.

Figure A.6: Monetary Policy Statement Surprises and Interest Rates

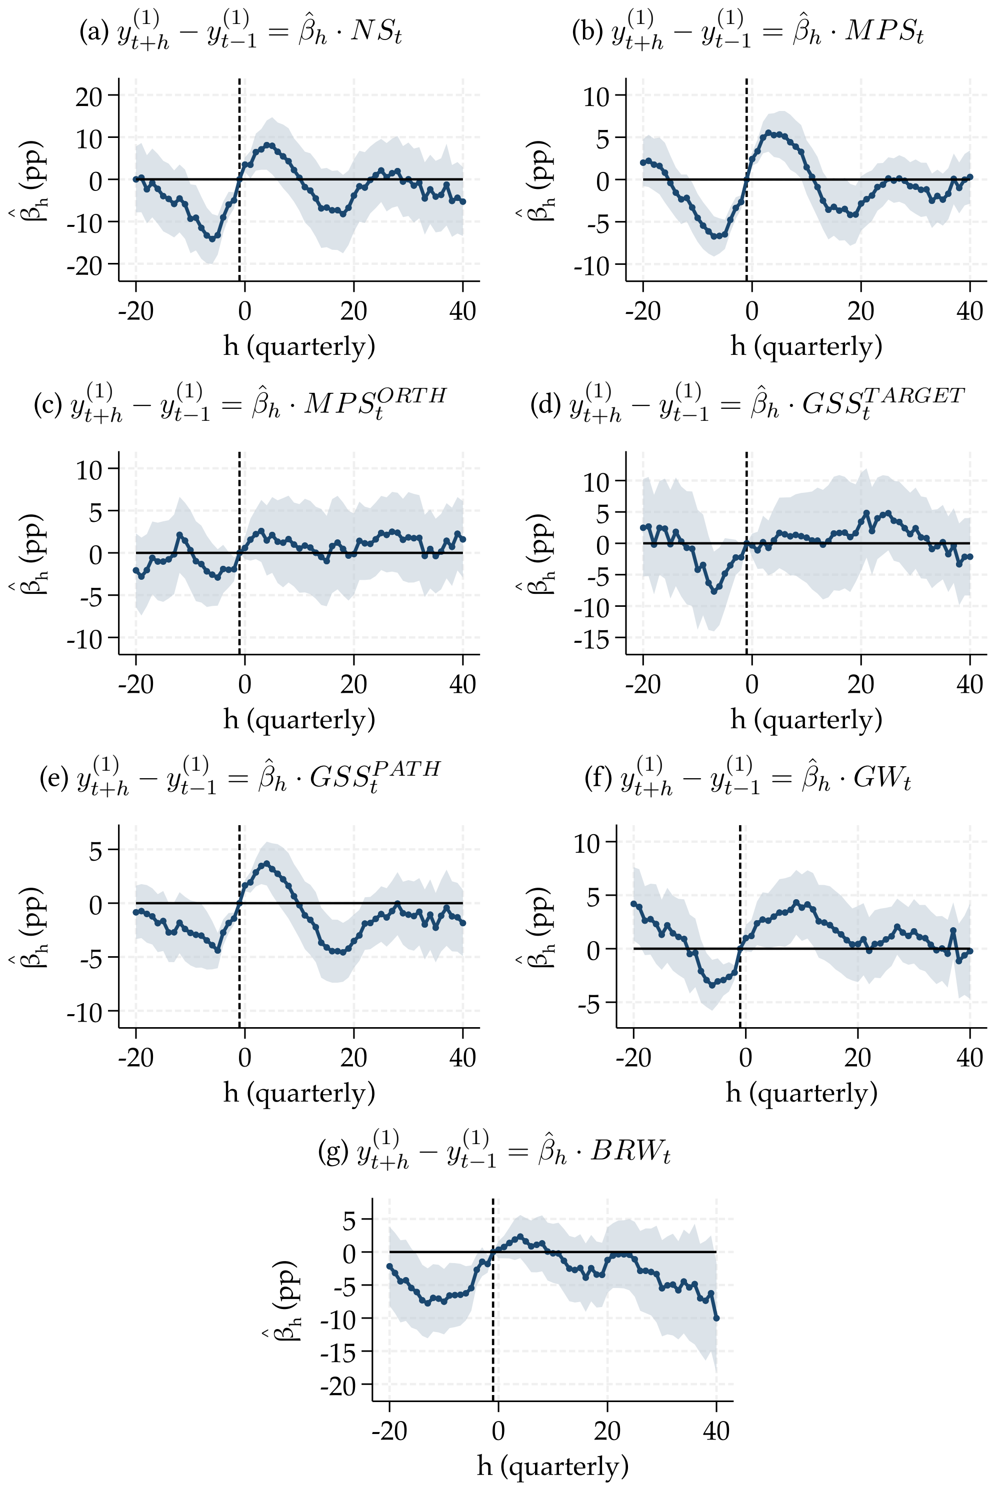

Note: The figure plots the sequence $\{\hat{\beta}_h\}_{h=-20}^{40}$ estimated from Equation (1). The outcome variable is the change in the one-year $(y^{(1)}_t)$ Treasury zero-coupon bond yield. All specifications are estimated at quarterly frequency using quarterly aggregated monetary policy surprises: NS from Nakamura and Steinsson [2018] updated by Acosta et al. [2024] (1995Q1–2023Q2); MPS and MPS^{ORTH} from Bauer and Swanson [2023] (1988Q1–2023Q4); GSS^{TARGET} and GSS^{PATH} from Gürkaynak et al. [2005] updated by Acosta et al. [2024] (1995Q1–2023Q2); GW from Gorodnichenko and Weber [2016] (1990Q1–2009Q4); BRW from Bu et al. [2021] (1994Q1–2020Q4). Confidence intervals are constructed at the 90 percent level using heteroskedasticity-robust standard errors.

Figure A.7: Monetary Policy Statement Surprises and Interest Rates: Long Daily Windows

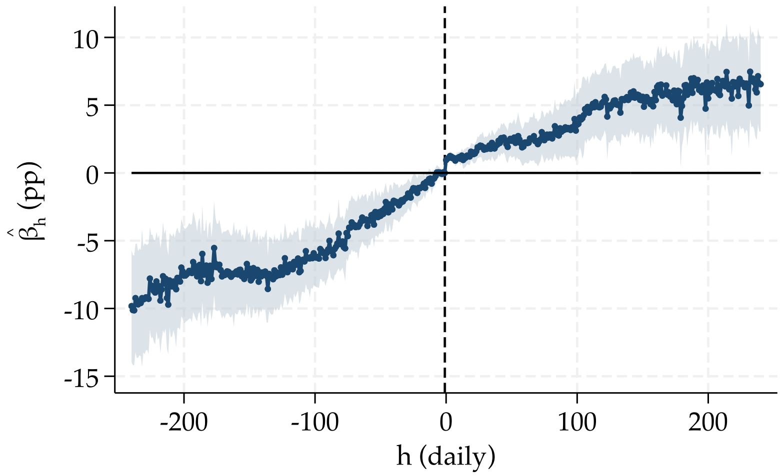

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-240}^{240}$ estimated from Equation (1) at the daily frequency. Monetary policy statement surprises come from Acosta et al. [2025]. The outcome variable is the change in the one-year $(y_t^{(1)})$ Treasury zero-coupon bond yield. Confidence intervals are constructed at the 90% level with heteroskedasticity robust standard errors.

Figure A.8: Stock Returns Around Monetary Policy Statement Surprises

(a) Baseline Sample (2002-2023)

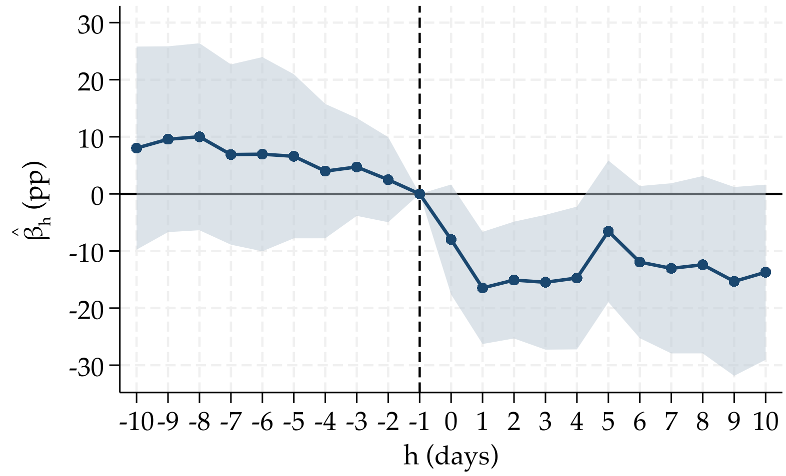

(b) Longer Sample (1994-2023)

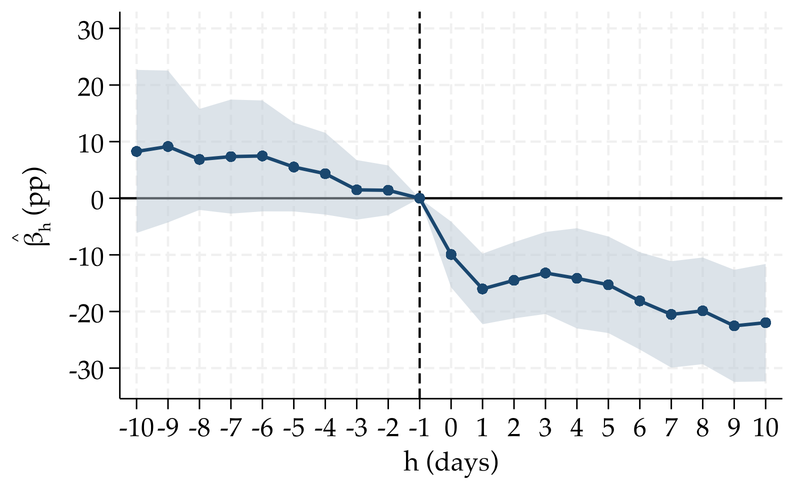

Note: This figure shows cumulative aggregate stock returns around FOMC meetings in response to monetary policy statement surprises. Returns are computed for a value-weighted portfolio of equities issued by firms in our sample. Panel (a) uses the baseline sample period 2002–2023, reflecting the availability of corporate bond data used elsewhere in the paper; Panel (b) extends the analysis to the full period for which monetary policy statement surprises are available, 1994–2023. For each meeting at date t, cumulative returns are measured relative to day $t-1$ and plotted from $t-10$ to $t+10$. Monetary policy statement surprises, $stmt_t$, come from Acosta et al. [2025] and are normalized to correspond to a one percentage point high-frequency increase in the one-year Treasury yield.

Figure A.9: Bond Returns Around Monetary Policy Statement Surprises

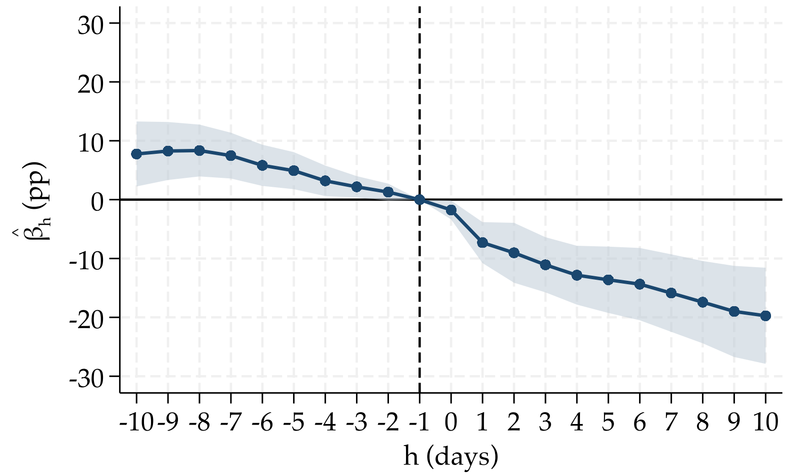

Note: This figure shows cumulative returns on the aggregate value-weighted portfolio of corporate bonds around FOMC meetings in response to monetary policy statement surprises. Bond returns are computed using TRACE transaction data by constructing bond-day prices as volume-weighted averages of clean trades. For each meeting at date t, cumulative returns are measured relative to day $t-1$ and plotted from $t-10$ to $t+10$. Monetary policy statement surprises, $stmt_t$, come from Acosta et al. [2025] and are normalized to correspond to a one percentage point high-frequency increase in the one-year Treasury yield.

Figure A.10: Return of Portfolio of Bonds and Stocks Around Monetary Policy *Event* Surprises

(a) Aggregate Portfolio of Bonds and Stocks

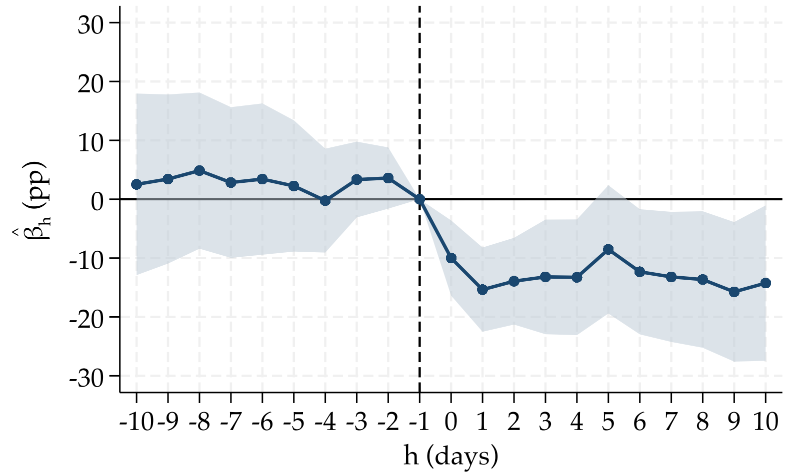

(b) Aggregate Portfolio of Stocks

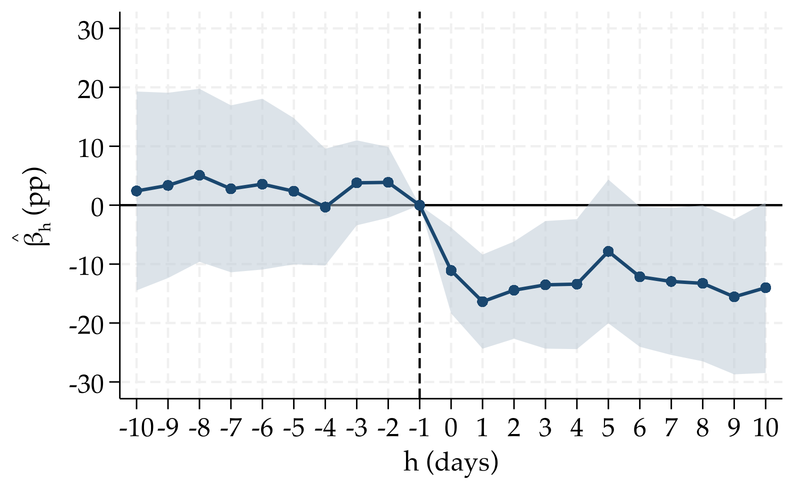

(c) Aggregate Portfolio of Bonds

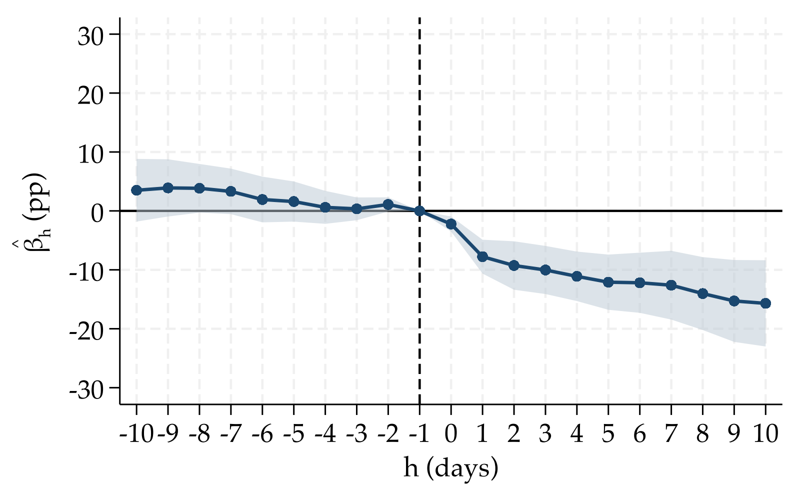

Note: This figure shows cumulative returns around FOMC meetings in response to monetary policy *event* surprises, me_t, from Acosta et al. [2025], for the sample 2002-2023. Panel (a) reports results for the aggregate value-weighted portfolio of bonds and equities, W; Panel (b) reports the corresponding value-weighted portfolio of equities only; Panel (c) reports the corresponding value-weighted portfolio of corporate bonds only. For each meeting at date t, cumulative returns are measured relative to day $t-1$ and plotted from $t-10$ to $t+10$. The event surprise captures policy-relevant information released over the full FOMC communication event, including the post-meeting press conference when applicable. Shaded areas denote 95% confidence intervals with standard errors clustered at the meeting level.

Figure A.11: Cross-Validation of Return Sensitivity Estimates: *Event* Surprises

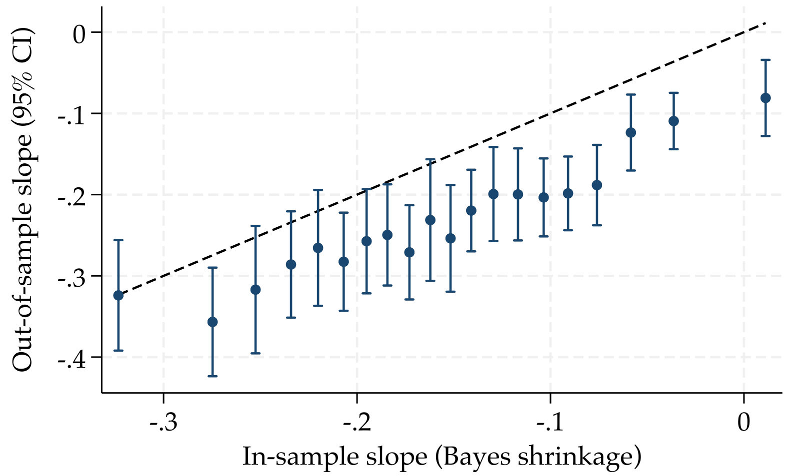

Note: This figure plots the out-of-sample return-shock slope estimated on held-out events (y-axis) against the average in-sample β_i^{Bayes} estimated on training events (x-axis) for 20 bins of firms. Error bars show 95% confidence intervals based on standard errors clustered by firm and event. The dashed line is the 45-degree line.

Figure A.12: Cross-Validation of Investment-Q Sensitivity Estimates: Characteristics-Based

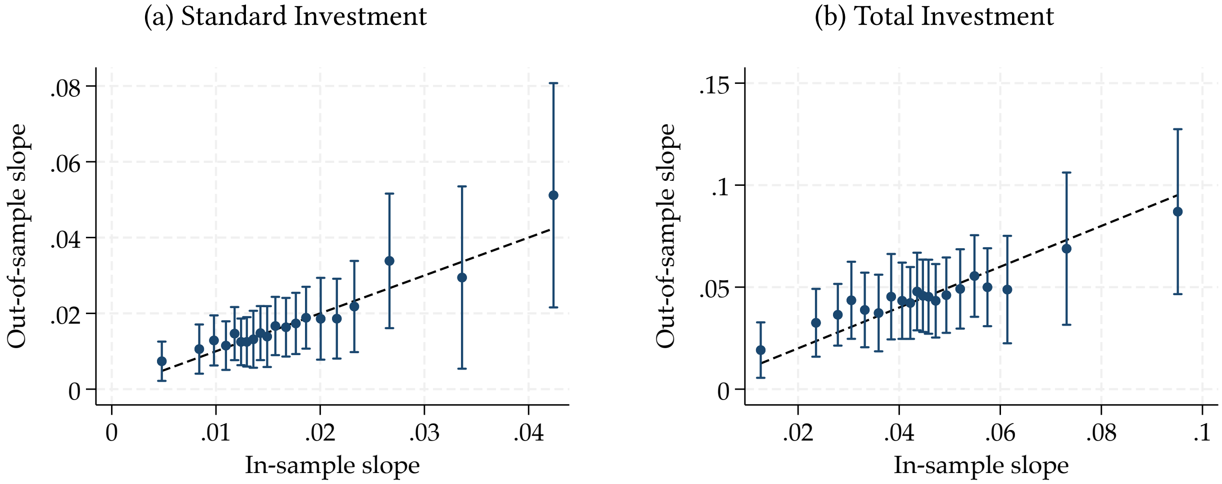

Note: Each panel plots the average out-of-sample investment-Q slope (y-axis) against the average in-sample $\gamma_{g(i)}^{Bayes}$ (x-axis) for 20 bins of firms. Error bars show 95% confidence intervals. The dashed line is the 45-degree line.

Figure A.13: Return Sensitivity and Investment-Q Sensitivity: Characteristics-Based

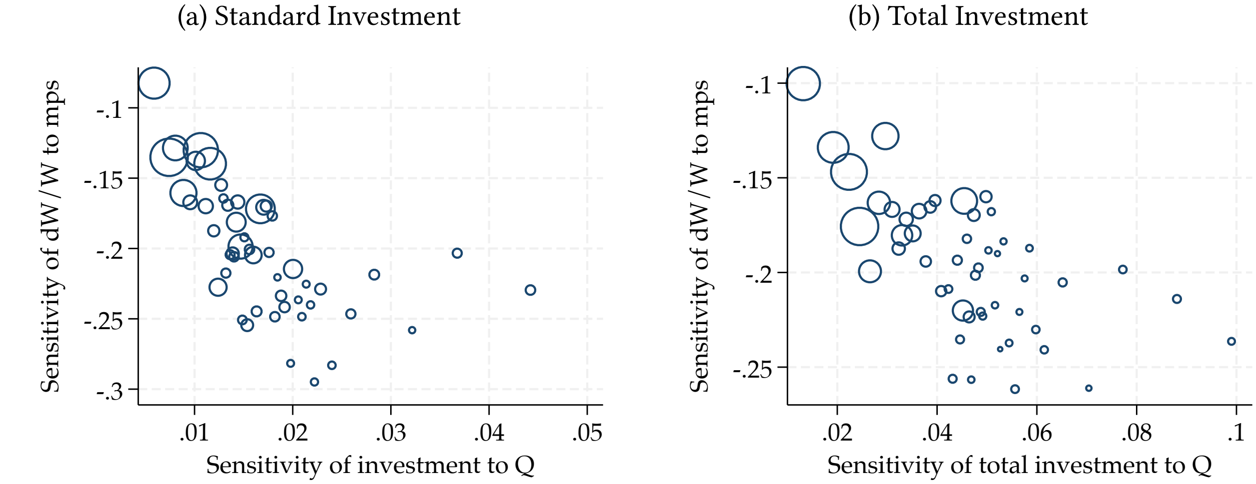

Note: Each panel plots value-weighted bin averages of return sensitivity to monetary policy (β_i^{Bayes}, y-axis) against investment-Q sensitivity ($\gamma_{g(i)}^{Bayes}$, x-axis). Firms are sorted into 50 bins based on $\gamma_{g(i)}^{Bayes}$. Point size reflects the total W-portfolio value of the firms in each bin.

D Additional Details on Q-Theory-Based Framework

D.1 Discussion of Assumption 1: Net Working Capital

Net working capital is defined as current assets minus current liabilities. In our setting, the main current assets are cash and short-term investments, accounts receivable, and inventories, while the main current liability (excluding short-term debt) is accounts payable. On the assets side, in our Compustat sample, about 33% of current assets are cash and short-term investments. We follow the corporate cash literature in treating Compustat cash and short-term investments (che) as a portfolio of cash and very short-maturity interest-bearing securities (e.g., deposits, commercial paper, and Treasury bills). Recent evidence shows that these short-term investments reprice quickly with money-market rates, so that che effectively behaves like a short-duration, floating-rate asset (e.g., Ysmailov [2021]). Accounts receivable function as short-maturity credit extended to customers and, in our sample, represent about 35% of current assets and roughly 64 days of sales. Given their rapid turnover, receivables are low-duration assets and reprice quickly as firms adjust credit terms to financial conditions. Inventories represent, in our sample, about 22% of current assets, and also have short effective duration because firms replenish them frequently: in our sample, Days Inventory Outstanding is also about two months. Since inventories tie up working capital whose opportunity cost is the short-term interest rate, the economic return to holding inventory adjusts quickly when interest rates change. On the liability side, accounts payable similarly function as short-term trade credit extended by suppliers. In our sample, they represent about 41% of current liabilities (excluding debt in current liabilities) and are typically due within 30–90 days, making them short-duration liabilities.[40] Other non-debt current liabilities primarily consist of accrued expenses, taxes payable, and deferred revenues. Because these items are typically settled within a few months and their implicit financing costs adjust with short-term interest rates, they also exhibit low effective duration. For these reasons, the main components of net working capital behave like low-duration assets and liabilities, implying that their market values should exhibit minimal sensitivity to high-frequency monetary policy surprises.

[40] In our sample, the average Days Payable Outstanding is about 67.