Pre-Trends Ahead of Monetary Policy Surprises: Recovering the Response of Corporate Investment[*]

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

July 11, 2026

[PRELIMINARY AND INCOMPLETE]

Abstract

How should high-frequency monetary policy surprises be used when the outcome of interest is only observed at a lower frequency? We first document a business-cycle endogeneity problem that arises when monetary policy surprises are aggregated to monthly or quarterly frequencies—surprise policy easings (resp. tightenings) tend to occur amid broader easing (resp. tightening) cycles. These low-frequency trends in pre-event Treasury yields are not reliably removed by standard controls without also weakening the relevance of the remaining policy variation. Motivated by this evidence, we next develop an alternative Q-theory-based approach for estimating the corporate investment response to monetary policy. Our approach combines high-frequency identification of firms' valuation responses to policy surprises with low-frequency firm-level investment-Q sensitivities. Our baseline estimate is: a one-unit monetary tightening, normalized to a one percentage point increase in the one-year Treasury yield, reduces the annual tangible investment of non-financial, non-utility public firms by 0.38% of GDP, with 24% of this effect driven by a novel firm-heterogeneity channel.

[*]This is a preliminary and incomplete draft. An earlier version of this project was circulated under the title "Monetary Policy and Corporate Investment: Sufficient Statistics with Heterogeneous Firms." We would like to thank Nicolas Crouzet, Janice Eberly, Samuel Hanson, Emi Nakamura, Dimitris Papanikolaou, Jon Steinsson, David Thesmar, Olivier Wang, and seminar participants at the SF Fed, UCLA, Chicago Booth, and Northwestern Kellogg for insightful discussion and suggestions.

[†]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 the economy? Despite being a fundamental question, the response of real macroeconomic outcomes to monetary policy remains challenging to pin down empirically. The difficulty, of course, is identification—monetary policy responds to economic conditions, meaning that policy actions are correlated with the same business-cycle forces that also shape subsequent macroeconomic outcomes.

The high-frequency identification approach has offered the literature a promising way to overcome this endogeneity problem by isolating unexpected monetary policy news using intraday interest-rate movements in narrow windows around policy announcements.[1] However, many outcomes of interest are only measured at the monthly or quarterly frequency (e.g., investment). For such outcomes, one common approach in the empirical literature is to aggregate high-frequency monetary policy surprises at the corresponding monthly or quarterly frequency, and then use the aggregated series as the event-time shock in an event-study analysis.[2] Although this is a natural extension of the high-frequency approach, it also changes the identifying assumption: when the outcome of interest is only observed at a monthly or quarterly frequency, the relevant shock must be orthogonal not only to news arriving within a narrow announcement window, but also to the broader business-cycle forces unfolding over the same monthly or quarterly horizon.

In the first part of this paper, we show that these sorts of lower-frequency monetary policy event studies suffer from a *business-cycle endogeneity* problem. In particular, we document that expansionary (resp. contractionary) policy surprises tend to occur during broader loosening (resp. tightening) cycles. In an event-study framework, this pattern manifests as pronounced low-frequency pre-trends in the one-year Treasury yield prior to a monetary policy surprise. Such pre-trends undermine the causal interpretation of low-frequency event-study designs, as they make it difficult to disentangle the effects of monetary policy from the effects of coincident business-cycle conditions.[3]

Though we have not developed an alternative empirical methodology that circumvents this business-cycle endogeneity problem generally, we do propose a new empirical approach for estimating how *corporate investment* responds to monetary policy. We present this alternative, Q-theory-based approach in the second part of this paper. Our 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 corporate-investment re-

[1] Studies using high-frequency identification include Kuttner [2001], Cochrane and Piazzesi [2002], Gürkaynak et al. [2005], Gertler and Karadi [2015], Gorodnichenko and Weber [2016], Nakamura and Steinsson [2018], Jarociński and Karadi [2020], and Bauer and Swanson [2023].

[2] See footnote 11 for a review of this literature.

[3] This lower-frequency finding echoes the higher-frequency predictability detailed in Bauer and Swanson [2023].

sponse into an average-firm channel and a firm-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 0.33% of U.S. GDP for the firms in our sample (or roughly 0.38% of GDP when extrapolated to all non-financial, non-utility public firms), with a meaningful share of this response attributable to firm-level heterogeneity.

Our first objective in this paper is to document an identification challenge that arises when high-frequency monetary policy surprises are taken outside of the narrow windows in which they are constructed. Note that the low-frequency usage of high-frequency surprises is common in the empirical literature, particularly when the outcome of interest is observed only monthly or quarterly (see footnote 11 for a review). While identification in a narrow intraday window requires that the FOMC communication be the dominant source of relevant new information in that narrow window, the maintained assumption in lower-frequency event studies is stronger: at a monthly or quarterly horizon, identification requires that the surprise be orthogonal to the broader macroeconomic forces evolving over that same horizon.

To investigate this lower-frequency orthogonality assumption, we implement a standard event-study design using the statement surprise series of Acosta et al. [2025].[4] Specifically, we aggregate meeting-level surprises to the quarterly frequency, and trace the path of the one-year Treasury yield around these quarterly-aggregated surprises. We find sizable pre-trends in the one-year Treasury yield prior to the quarter of the monetary policy surprise, with surprise easings tending to occur after a period of falling yields (and surprise tightenings after rising yields). This result is not a mechanical artifact of quarterly aggregation. Similar pre-trends appear with monthly-aggregated surprises, and at the daily frequency when the event study is run with a wide event window. In short, low-frequency monetary policy event studies suffer from a business-cycle endogeneity problem, since measured monetary policy surprises are correlated with slow-moving interest-rate and business-cycle dynamics that begin well before the policy surprise itself.[5]

We next show that this business-cycle endogeneity concern is not easily resolved by adding controls. We consider standard macro-financial controls,[6] and we also control directly for lags of quarterly yield changes. In all cases, we find that controls either do not eliminate low-frequency pre-trends, or controls absorb too much of the event-time policy variation, or both. Moreover, we also demonstrate an important practical implication of this business-cycle endogeneity problem—the estimated response of macroeconomic outcomes such as unemployment, inflation, GDP, and

[4]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.

[5]To be clear, this is an in-sample finding. For example, as discussed in Bauer and Swanson [2023], the pre-trends we identify could be consistent with market participants having incomplete information about the Fed's policy rule. For purposes of demonstrating that business-cycle endogeneity is a concern, however, the interpretation is secondary.

[6]Specifically: (i) we control for four lags of GDP growth, unemployment, and inflation, and (ii) we include the controls used in the Bauer and Swanson [2023] orthogonalization procedure.

investment can be highly sensitive to the specific controls used. Thus, the conclusions one draws about monetary policy's effects are likely to be shaped by the chosen control set rather than being robustly identified by the event-study variation itself. Oster [2019] diagnostics similarly imply wide ranges of plausible effects.

The business-cycle endogeneity problem that we document is also not specific to our baseline surprise measure. We repeat the analysis across a broad set of commonly used high-frequency monetary policy surprise series, including those from Acosta et al. [2025], Bauer and Swanson [2023], Bu et al. [2021], Gorodnichenko and Weber [2016], Gürkaynak et al. [2005], Miranda-Agrippino and Ricco [2021], and Nakamura and Steinsson [2018]. Across these surprise series, quarterly-aggregated surprises exhibit significant pre-trends, weak relevance for contemporaneous one-year-yield movements, or both.

Finally, it is important to clarify the scope of this business-cycle endogeneity problem. In short, we do not argue that high-frequency identification is invalid for high-frequency outcomes. In fact, when we estimate the response of asset prices in narrow windows around FOMC announcements—including the one-year yield, the S&P 500, the 10-year yield, and the 10-year real yield—the estimated effects are stable across the same control sets that generate substantial sensitivity at the quarterly frequency. The available evidence thus points to a horizon-specific identification problem. High-frequency monetary policy surprises remain useful for estimating the immediate asset-price effects of policy news, but difficulties arise when those same surprises are used as if they were exogenous shocks to lower-frequency macroeconomic outcomes.

Our second objective in this paper is to develop an alternative approach to estimating the corporate investment response to monetary policy. Because corporate investment is not observed at high frequency—and in light of the business-cycle endogeneity problem just discussed—we do not try to estimate this response by directly projecting quarterly investment on quarterly-aggregated monetary policy surprises. Instead, our alternative approach combines two ingredients. First, we preserve the high-frequency identifying variation to estimate how firms' market values respond to monetary policy news within a narrow window around FOMC 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. Correspondingly, changes in the firm's enterprise value map into changes in investment through the firm's investment-Q sensitivity. We can therefore 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 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. 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%.[7] One additional advantage of our Q-theory-based framework is its transparency: it provides a first-principles mapping from high-frequency valuation responses to lower-frequency investment responses, and makes explicit the objects needing to be estimated empirically. While one could alternatively use high-frequency-identified asset-price responses as moments to discipline a richer structural model of firms' investment and financing decisions—and indeed we view this as a promising direction for future research—the identification issues documented above make a simple, transparent, sufficient-statistics benchmark the natural first step.

Detailing our Q-theory-based framework, we first need a market-based measure of firms' enterprise values (where market prices are needed to capture changes in firm value within tight announcement windows). 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.[8] 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 this approximation, the change in firm i's enterprise value following a monetary policy shock dr_t corresponds to the change in its combined equity-and-bond portfolio value, dW_{it}. We let β_i denote the high-frequency return sensitivity of firm i's W portfolio to monetary policy shocks (i.e., $\beta_i = \frac{dW_{it}/dr_t}{W_{it}}$). Q-theory then maps these high-frequency valuation changes into investment changes through firm i's investment-Q sensitivity, γ_i. Together, these two objects yield a simple firm-level sufficient statistic: firm i's investment response to a monetary policy shock is governed by the product $\gamma_i \times \beta_i$.

Aggregating across firms and scaling by the aggregate portfolio value W_t delivers a tractable decomposition characterizing the aggregate investment response to a monetary policy shock:

$$\frac{dI_t/dr_t}{W_t} = \mathbb{E}_w[\gamma_i] \times \mathbb{E}_w[\beta_i] + \mathbb{C}\text{ov}_w(\gamma_i, \beta_i),$$

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

[8] 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.

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 induces an average valuation response ($\mathbb{E}_w[\beta_i]$), which then translates into investment through the average investment-Q sensitivity ($\mathbb{E}_w[\gamma_i]$). The covariance term captures the role of firm-level 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.

We then take this sufficient-statistics framework to the data. Doing so requires two sets of inputs, which are estimated in separate steps. First, we estimate the high-frequency return sensitivity of firm-level W portfolios to monetary policy surprises, which provides the β_i inputs. Second, we estimate how firm-level investment responds to changes in valuation, which provides the γ_i inputs.

We start by estimating firm-level return sensitivities to monetary policy surprises (β_i). We construct firm-level W portfolios using equity prices from CRSP and corporate bond prices from TRACE. Monetary policy surprises are again taken from Acosta et al. [2025]. For each firm, we estimate β_i by regressing the two-day cumulative return on its W portfolio around FOMC announcements on the monetary policy statement surprise.

The average return sensitivity to a one-unit monetary policy surprise is -0.16. This indicates that monetary policy news produces sharp and economically meaningful movements in firms' market valuations. Because firm-specific estimates are likely to be noisy, we also verify that the cross-sectional heterogeneity in return sensitivities we identify is not merely an artifact of sampling error by conducting a 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 is firm-level investment-Q sensitivities (γ_i). Using annual Compustat data (1994–2023, excluding 2020), we run standard investment-Q regressions for both standard (tangible) investment and for a broader "total investment" measure that incorporates intangibles [Eisfeldt and Papanikolaou, 2013; Peters and Taylor, 2017].[9] Again, because our firm-specific estimates are likely to be noisy, we conduct a 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 sensi-

[9]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.

tive to monetary policy surprises. That is, the covariance term in the decomposition above is negative, $\mathbb{C}\text{ov}_w(\gamma_i, \beta_i) < 0$. Economically, this negative covariance implies that surprise tightenings disproportionately reduce valuations for precisely the firms whose investment responds more strongly to such valuation changes. This firm-heterogeneity channel therefore *amplifies* the aggregate investment response relative to what an average-firm-only calculation would imply.

Finally, we quantify the response of corporate investment to monetary policy. Plugging our baseline estimates into the sufficient-statistics decomposition above, we find that a one-unit surprise tightening reduces (tangible) investment by about 0.36% of the value of the aggregate W portfolio. Since the aggregate W portfolio is roughly 0.9-times U.S. GDP in our sample, this corresponds to an annual investment reduction of approximately 0.33% of U.S. GDP among the firms we study.[10] 24% of this 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 endogeneity problem faced by low-frequency monetary policy event studies. Section 3 lays out our alternative, Q-theory-based framework for estimating the corporate investment response to monetary policy. Section 4 takes the framework to the data and quantifies the implied corporate investment response. Section 5 concludes.

2 Business-Cycle Variation in Monetary Policy Surprises

Since many macroeconomic variables of interest—such as GDP, inflation, and investment—are measured at monthly or quarterly frequencies, a large empirical literature aggregates high-frequency monetary policy surprises to the corresponding lower frequency and then uses these aggregated surprises to estimate the dynamic effects of monetary policy on such outcomes.[11] In doing so, this approach implicitly assumes that the monthly- or quarterly-aggregated surprises are orthogonal to broader macroeconomic dynamics unfolding over the same horizon.

In this section, we call that identifying assumption into question by highlighting a *business-cycle endogeneity* problem. Specifically, we document that surprise easings (resp. tightenings) are

[10] The investment reduction is 0.38% if the same per-dollar response is extrapolated to the remaining non-financial, non-utility public firms that are not included in our sample. See Section 4.1 for sample construction details.

[11] For papers studying the response of corporate investment (our focus in this paper) to monetary surprises at monthly or quarterly frequencies, see e.g., Ottonello and Winberry [2020], Caglio et al. [2021], Crouzet [2021], Cloyne et al. [2023], Döttling and Ratnovski [2023], Ferreira et al. [2023], Hsu et al. [2023], Jeenas and Lagos [2024], Jungherr et al. [2024], Gnewuch and Zhang [2025], Kroen et al. [2025], Beyhaghi et al. [2026], Drechsel et al. [2026], Jeenas [2026], Lakdawala and Moreland [2026], Perez-Orive et al. [2026], and Selgrad and Siani [2026]. For papers studying other low-frequency outcomes, see also Cochrane and Piazzesi [2002], Faust et al. [2003], Faust et al. [2004], Barakchian and Crowe [2013], Gertler and Karadi [2015], Ramey [2016], Stock and Watson [2018], Caldara and Herbst [2019], Kalemli-Özcan [2019], Lakdawala [2019], Jarociński and Karadi [2020], Lunsford [2020], Miranda-Agrippino and Rey [2020], Paul [2020], Bu et al. [2021], Miranda-Agrippino and Ricco [2021], Nunes et al. [2022], Bauer and Swanson [2023], Bilbiie and Känzig [2023], Bauer et al. [2024], Damast et al. [2025], Abramson et al. [2026], Gamber [2026], Loualiche et al. [2026], and Pazarbaşı [2026].

systematically embedded in broader easing (resp. tightening) cycles. In an event-study framework, this business-cycle endogeneity appears as pronounced pre-trends in macroeconomic variables prior to a monetary policy surprise. These pre-trends undermine the causal interpretation of empirical designs that directly map high-frequency monetary policy surprises into lower-frequency outcomes, since estimated responses may partly reflect the continuation of pre-existing macroeconomic conditions rather than the causal effect of the surprise alone.

The remainder of this section examines this business-cycle endogeneity problem in further detail. This problem, in turn, points to the need for alternative empirical methodologies for estimating the effects of monetary policy on lower-frequency outcomes. Starting in Section 3, we develop and implement one such methodology in the context of corporate investment.

2.1 Low-Frequency Pre-Trends in Monetary Policy Event Studies

Our baseline measure of monetary policy news is the statement surprise series ($stmt_t$) 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.[12] The resulting factor is rescaled so that a unit change corresponds to a one percentage 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.[13]

Main Result. Figure 1 presents event-study estimates of the effect of monetary policy statement surprises on the one-year Treasury yield at both the daily and quarterly frequency.[14] We plot the estimated coefficients $\hat{\beta}_h$ from the following OLS regression in event time:

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

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

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

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

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 $\{\hat{\beta}_h\}^{20}_{h=-20}$ traces the event-time impulse-response of the one-year Treasury yield around a monetary policy statement surprise.[15] 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 restrict attention to scheduled FOMC statement surprises, and also 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 242 monetary policy statement surprises.

Panel (a) of Figure 1 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 102 bps increase in the one-year Treasury yield on the day of the surprise, and an overall 165 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 recent macroeconomic and financial news, which move yields prior to FOMC meetings, predict the magnitude and sign of monetary policy surprises.[16]

Panel (b) then 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.[17] β_h in Equation (1) then captures the response of the one-year Treasury yield h quarters after the quarterly-aggregated surprise. Importantly, Panel (b) reveals sizable pre-trends in the one-year Treasury yield prior to a quarterly statement surprise.[18] These pre-trends imply that quarterly-aggregated statement surprises are not isolated from broader pre-existing interest-rate and business-cycle dynamics.

We focus on quarterly aggregation here since the response of investment to monetary policy—our main focus in this paper—is typically studied at quarterly frequencies. However, the precise level of aggregation is not critical. Low-frequency pre-trends can be seen for monthly-aggregated monetary policy surprises (see Appendix Figure A.2), or even for daily monetary policy surprises

[15] We restrict the sample to announcement periods only, but results are essentially unchanged if we set $stmt_t = 0$ in periods without an FOMC announcement.

[16] Several papers show that monetary policy surprises are predictable from information observed right before the FOMC announcement: Greenbook forecasts and forecast revisions based on the most recent projections available before the announcement [Miranda-Agrippino, 2016; Miranda-Agrippino and Ricco, 2021]; 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]; option-implied Treasury skewness averaged over the month preceding the meeting [Bauer and Chernov, 2024]; and three-month pre-meeting changes in stock prices, the yield curve slope, and commodity prices [Bauer and Swanson, 2023].

[17] Appendix Figure A.3 shows that results are insensitive to whether we sum statement surprises within a quarter or weight surprises by their timing within the quarter.

[18] Appendix Figure A.1 quantifies this by forecasting quarterly statement surprises with lags of quarterly yield changes: one lag forecasts $stmt_t$ with an R^2 of 26%, while ten lags deliver an R^2 of 33%.

when the window is extended to 240 trading days (roughly a year) prior to the announcement (see Figure 2 below). In other words, monetary policy surprises are correlated with slow-moving business-cycle conditions that start to unfold long before the surprise itself.[19]

The pre-trends in Panel (b) of Figure 1 can be understood as extending, 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 to longer-term macroeconomic and financial news.[20]

Regardless of interpretation, Figure 1 makes clear that quarterly-aggregated monetary policy surprises are systematically related to past yield changes, and hence to prior business-cycle dynamics. This pattern raises a general concern for empirical designs that aggregate high-frequency monetary policy surprises to match lower-frequency macroeconomic outcomes (such as quarterly investment). Specifically, the pre-trends that we document here imply that these empirical designs may fail to isolate the causal effects of monetary policy, with estimated responses instead partly reflecting the continuation of pre-existing business-cycle trends.

Wide-Window Daily Event Study. As mentioned above, the business-cycle endogeneity problem highlighted by Figure 1 is not driven mechanically by time aggregation to the quarterly level. To show this, Figure 2 provides an event study around *daily* monetary policy statement surprises using a wider window of 240 days. Figure 2 exhibits similar low-frequency pre-trends to those in our baseline quarterly event study in Figure 1.

Clarifying Scope of Business-Cycle Endogeneity Problem. Before continuing, it is important to be precise about the scope of the business-cycle endogeneity problem. We *do not claim* that the low-frequency pre-trends documented here will necessarily undermine identification in high-frequency settings. Indeed, the Q-theory-based methodology for estimating the corporate investment response to monetary policy that we propose in Sections 3 and 4 will rely on estimating the high-frequency response of asset prices to monetary policy surprises. Moreover, we further justify the validity of such high-frequency designs in Section 2.4 below.

Our critique is instead directed at lower-frequency applications: once monetary policy surprises are used at monthly or quarterly frequencies, they are unlikely to provide a clean source

[19] We highlight one other curious feature in Panel (b) of Figure 1, which is that the β_0 coefficient is roughly 4. That is, the one-year yield moves by much more than the surprise itself within the quarter of the surprise.

[20] Such an interpretation would constitute a rejection of full-information rational expectations, although it remains potentially consistent with market efficiency.

Figure 1: Path of One-Year Yield Around Monetary Policy Statement Surprise

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

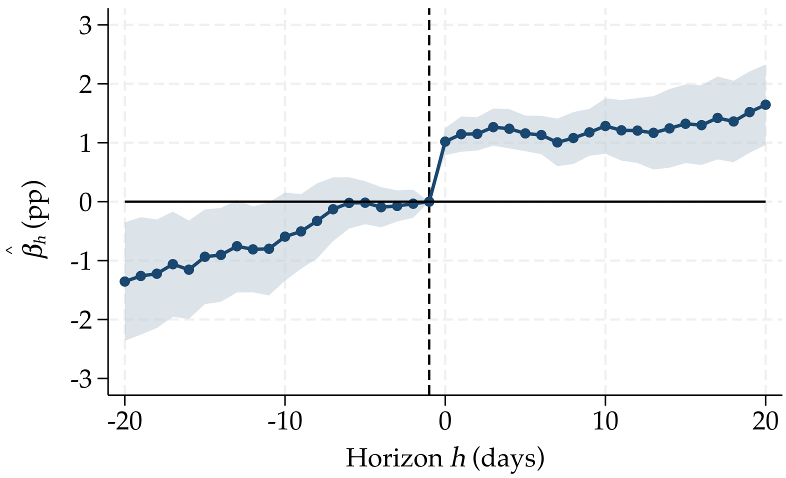

(b) $y^{(1)}_{t+h} - y^{(1)}_{t-1} = \hat{\beta}_h \cdot stmt_t$, Quarterly Frequency

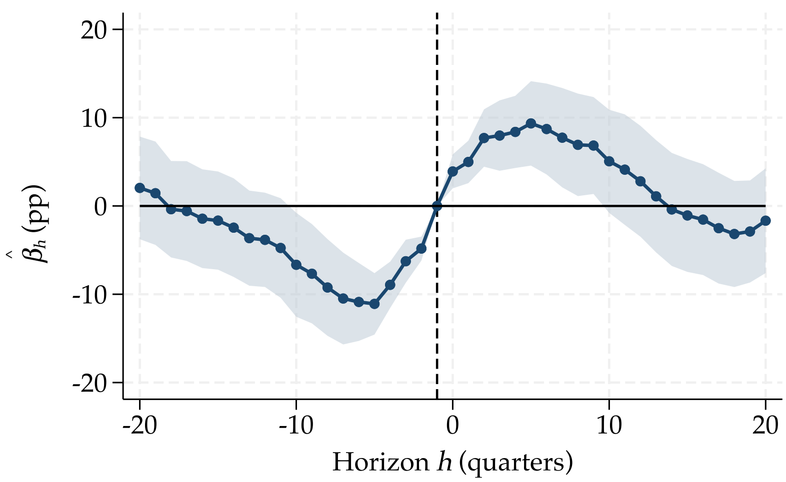

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-20}^{20}$ estimated from Equation (1). Monetary policy statement surprises come from Acosta et al. [2025]. The outcome variable is the change in the one-year ($y^{(1)}_t$) Treasury zero-coupon bond yield. Confidence intervals are constructed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

of variation for identifying the causal effects of monetary policy because they systematically co-move with broader business-cycle conditions.

The distinction is largely one of horizons. A high-frequency analysis of monetary policy surprises relies on the assumption that, within a narrow window around an FOMC statement, the monetary policy surprise itself is the dominant source of relevant new information. By contrast, using quarterly-aggregated surprises to identify the response of a quarterly outcome requires a stronger assumption: because the outcome-window extends over a much longer horizon, the

Figure 2: Path of One-Year Yield Around Monetary Policy Statement Surprise: Wide Daily Window

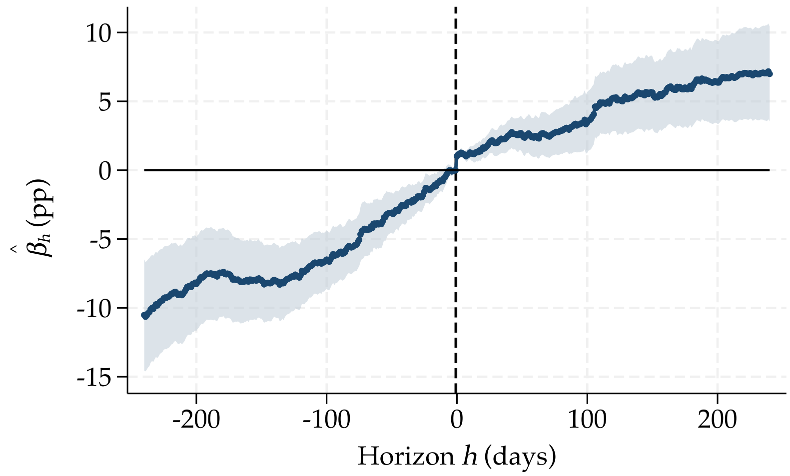

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-240}^{240}$ estimated from Equation (1) at the daily frequency. Monetary policy statement surprises come from Acosta et al. [2025]. The outcome variable is the change in the one-year $(y_t^{(1)})$ Treasury zero-coupon bond yield. Confidence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

aggregated surprise must be orthogonal to slow-moving macroeconomic forces that also affect the outcome over the same horizon. The pre-trends highlighted above indicate that this stronger assumption is unlikely to hold.[21]

2.2 Adding Controls Exacerbates Relevance and Sensitivity Concerns in Quarterly Event Studies

One solution to the business-cycle endogeneity problem just highlighted could, potentially, be to control for pre-existing macroeconomic conditions that are correlated with monetary policy surprises. The identifying assumption implicit in this approach is that the effect of monetary policy surprises on subsequent outcomes is conditionally unconfounded—that is, monetary policy surprises are as good as random after controlling for various business-cycle variables.

As is often the case, however, controlling for business-cycle variables is unlikely to fully restore exogeneity. Indeed, Figure 3 below illustrates that even with controls, quarterly-frequency specifications are likely to suffer from at least one of two central issues: controlling for standard macro variables may not remove pre-trends, and if it does, then the residualized monetary policy surprise that remains may no longer provide a relevant source of identifying variation. Moreover, Figure 4 highlights an additional issue: the estimated response of macroeconomic outcomes can

[21] Even if monetary policy surprises are only predictable ex post, their systematic relationship with prior yield movements nonetheless suggests that they are not orthogonal to broader business-cycle conditions.

be highly sensitive to the specific controls used. We detail these three issues in turn below.

Quarterly Monetary Policy Event Study with Controls: One-Year Yield. To begin, Figure 3 replicates the quarterly event study of Figure 1 but with the addition of four different sets of controls. In Panel (a) we utilize two sets of conventional business-cycle controls: (i) we control for the set of pre-announcement macroeconomic and financial variables used by Bauer and Swanson [2023] when constructing their "orthogonalized monetary policy surprise" measure,[22] and (ii) we control for four quarterly lags of real GDP growth, unemployment, and inflation.[23] In Panel (b), we directly control for: (i) the past eight lags of quarterly yield changes, and (ii) the past eight lags of quarterly yield changes plus all of the controls used in Panel (a).

Starting with the Bauer and Swanson [2023] controls in Panel (a) of Figure 3 (dashed red line), these controls only partially ameliorate pre-trends. This implies that the lower-frequency dynamics we uncover are not eliminated by their orthogonalization procedure. Moreover, the residualized monetary policy surprise is now only weakly related to the event-quarter change in the one-year yield ($h = 0$). The macro controls in Panel (a) (dashed green line)—which include four lags of real GDP growth, unemployment, and inflation—likewise fail to remove pre-trends. Looking next at Panel (b), directly controlling for eight lags of yield changes (dashed yellow line) or including all business-cycle and yield controls (dotted purple line) does mechanically remove pre-trends in yields.[24] However, these controls again appear to absorb much of the relevant event-time policy variation (the estimated coefficient even changes sign when all controls are included).

Overall, Figure 3 clearly illustrates the first two issues mentioned at the start of this subsection: namely, that controlling for standard business-cycle variables may not remove pre-trends, and if it does then any monetary policy surprise that remains may not be relevant. Although Figure 3 only provides a visual analysis of these two issues while looking specifically at the statement surprise series of Acosta et al. [2025], we will soon show in Section 2.3 that these issues exist across a variety of monetary policy surprise measures that are common in the literature.

Quarterly Monetary Policy Event Study with Controls: Macro Outcomes. Next, Figure 4 digs into the third issue with adding controls, which is that the estimated response of macroeconomic outcomes can be highly sensitive to the specific controls that are used. Figure 4 presents impulse-responses around quarterly-aggregated monetary policy surprises for four different macroeconomic outcomes: unemployment, inflation, GDP, and investment. The key takeaway from Figure 4 is that the estimated response of each of these outcomes is highly sensi-

[22]These variables include the nonfarm payrolls surprise (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 futures.

[23]These business-cycle controls are used in Ottonello and Winberry [2020], among others.

[24]Nonetheless, pre-trends in other variables still remain—see in particular Table 3 in Section 2.3.

Figure 3: Path of One-Year Yield Around Monetary Policy Statement Surprise: Adding Controls

(a) Business-Cycle Controls

(b) Direct Yield Controls

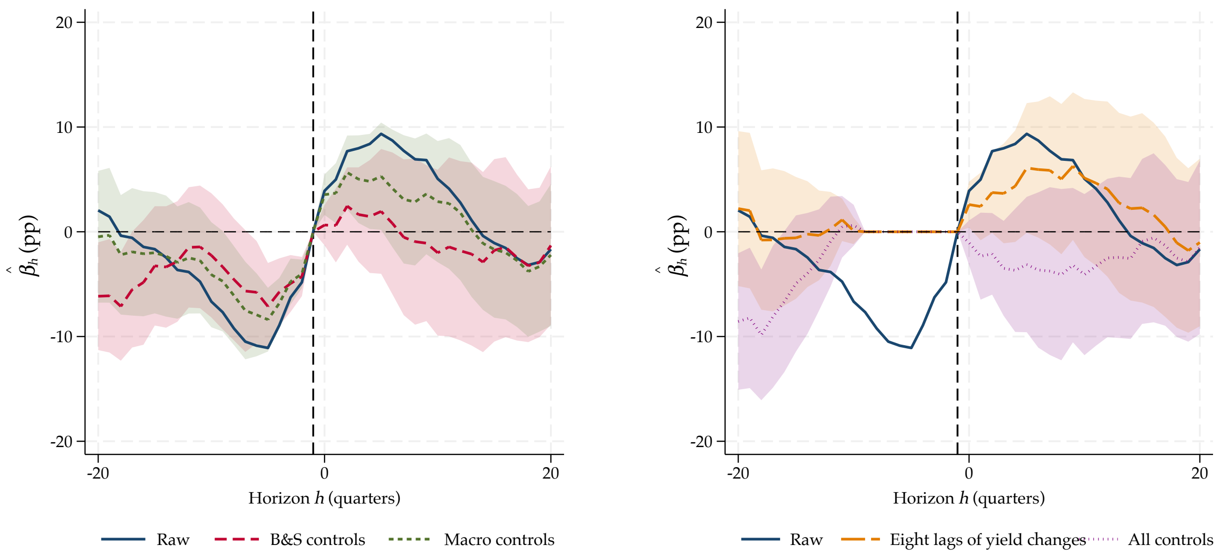

Note: This figure plots the sequence $\{\hat{\beta}^v_h\}^{20}_{h=-20}$ estimated from the specification

$$y^{(1)}_{t+h} - y^{(1)}_{t-1} = \alpha^v_h + \beta^v_h \cdot stmt_t + \gamma^v_h X^v_t + \varepsilon^v_{t,h},$$

where $stmt_t$ is the quarterly-aggregated statement surprise from Acosta et al. [2025] and X^v is a set of controls. Each line corresponds to a different set of variables v. B&S controls are the controls used by Bauer and Swanson [2023]: the most recent nonfarm payrolls surprise, employment growth, stock-market returns, yield-curve slope changes, commodity-price changes, and option-implied Treasury skewness. Macro controls are four lags of real GDP growth, inflation, and unemployment. Lags of yield changes are quarterly changes in the one-year yield in the eight quarters preceding the surprise. Confidence intervals are constructed at the 90% level with heteroskedasticity-robust standard errors. For visual clarity, we omit confidence intervals for the raw sequence (see Panel (b) of Figure 1).

tive to the specific controls used. This suggests that the conclusions one draws about monetary transmission depend materially on the chosen control set, rather than emerging robustly from the event-study variation itself.

One way to quantify the sensitivity of these estimates to omitted-variable bias is to construct bounds on plausible effect sizes using the methodology from Oster [2019].[25] Table 1 reports the implied range of effects of monetary surprises on unemployment, inflation, GDP, and investment at a one-year horizon under alternative assumptions about selection on unobservables. Even modest selection on unobservables relative to observed controls generates wide intervals. For example, if selection on other omitted variables is no stronger than selection on the path of prior

[25]This approach uses the change in the measured effect of the monetary surprise on the outcome when controls are added, along with the share of variation in the outcome explained by the shock and controls, to estimate the extent of selection on the controls. This can be used to gauge the range of plausible effects of the shock on the outcome under alternative assumptions about how correlated the shock is with unobservables not included as controls (i.e., other omitted variables).

yield changes, the effect of a one-unit contractionary monetary surprise on real GDP after one year ranges from -0.6% to 4.7%.[26] Moreover, even the Oster [2019] ranges change materially with the set of controls that we use.

We end by noting that this sort of fragility is, perhaps, unsurprising. Simply put, adding controls is unlikely to address endogeneity in a satisfactory way when the source of that endogeneity is not fully known. In our setting, once one recognizes that monetary policy "shocks" systematically comove with business-cycle conditions, adding controls can only remove the components of business-cycle variation captured by those specific controls. Because we do not know the precise sources of business-cycle comovement for monetary policy surprises, different control sets remove different sources of variation and therefore generate different estimated responses.

Table 1: Sensitivity of Estimated Effects on Low-Frequency Macro Outcomes: Oster [2019] Bounds

			Effect of monetary surprise on outcome at $t+4$			
			Controls for 8 lags of yield changes		All controls	
Outcome	$\hat{\beta}_{\text{short}}$	R^2_{short}	R^2_{long}	Oster [2019] bounds	R^2_{long}	Oster [2019] bounds
Unemployment Rate	-7.8	0.081	0.197	[-5.5, -0.6]	0.623	[0.4, 9.1]
Log CPI	7.0	0.044	0.109	[1.0, 5.1]	0.573	[-8.2, -0.7]
Log Real GDP	7.9	0.044	0.133	[-0.6, 4.7]	0.530	[-16.6, -4.5]
Log Net Private Investment	132.6	0.033	0.147	[2.6, 74.9]	0.524	[-208.7, -43.2]

Note: The columns $\hat{\beta}_{\text{short}}$ and R^2_{short} report results from $x_{t+4} - x_{t-1} = \beta_{\text{short}} \cdot stmt_t + \varepsilon_t$ for each outcome x. The two pairs of long columns differ in their control set: the first ("8 lags of yield changes") adds only the eight quarterly yield-change lags $\{\Delta y^{(1)}_{t-k}\}_{k=1}^{8}$; the second ("All controls") adds those plus four quarterly lags of GDP, unemployment, and inflation, and the B&S controls. The Oster [2019] range is calculated by varying $\delta \in [-1, 1]$ using $\beta^* = \beta_{\text{long}} - \delta\left(\beta_{\text{short}} - \beta_{\text{long}}\right) \frac{R^2_{\text{max}} - R^2_{\text{long}}}{R^2_{\text{long}} - R^2_{\text{short}}}$, where R^2_{long} is the R-squared of the regression with controls and $R^2_{\text{max}} = \min\{1, 1.3 \cdot R^2_{\text{long}}\}$.

2.3 Pre-Trends and Relevance Concerns Across Quarterly-Aggregated Monetary Policy Surprise Measures

While we have thus far focused on the statement surprise series of Acosta et al. [2025]—since this is our baseline measure of monetary policy news—we now document that the business-cycle endogeneity problem applies across a range of alternative monetary policy surprise measures that are commonly used in the literature.

[26] Recent work by Masten and Poirier [2026] shows that the Oster [2019] methodology can understate the likelihood that a coefficient changes sign. In this sense, even the wide intervals in Table 1 may understate the true uncertainty in the estimated effects of monetary surprises on these outcomes.

Figure 4: Sensitivity of Macro Impulse-Responses to Controls Used

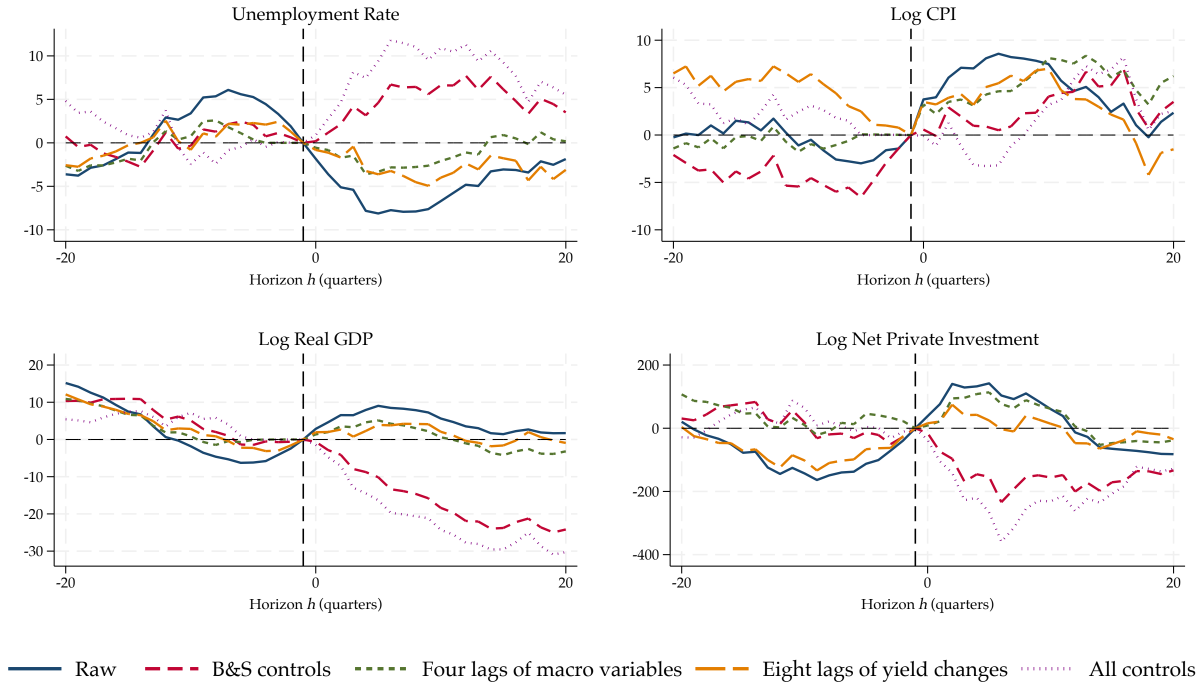

Note: This figure plots the sequence $\{\hat{\beta}_h^v\}_{h=-20}^{20}$ estimated from the specification

$$x_{t+h} - x_{t-1} = \alpha_h^v + \beta_h^v \cdot stmt_t + \gamma_h^v X_t^v + \varepsilon_{t,h}^v,$$

where x_t is an outcome x in quarter t, $stmt_t$ is the quarterly-aggregated statement surprise from Acosta et al. [2025], and X^v is a set of controls. Each line corresponds to a different set of variables v. B&S controls are the controls used by Bauer and Swanson [2023]: the most recent nonfarm payrolls surprise, employment growth, stock-market returns, yield-curve slope changes, commodity-price changes, and option-implied Treasury skewness. Macro controls are four lags of real GDP growth, inflation, and unemployment. Lags of yield changes are quarterly changes in the one-year yield in the eight quarters preceding the surprise.

We start by examining a variety of "off-the-shelf" measures of quarterly-aggregated monetary policy surprises. These include both the statement surprise series and the event surprise series from Acosta et al. [2025] (raw and orthogonalized),[27] as well as monetary policy surprise series from Bauer and Swanson [2023] (raw and orthogonalized), Bu et al. [2021], Gorodnichenko and Weber [2016], Gürkaynak et al. [2005] (target and path; updated by Acosta et al. [2024]), Miranda-Agrippino and Ricco [2021] (orthogonalized), Nakamura and Steinsson [2018] (updated by Acosta et al. [2024]), and Romer and Romer [2004] (updated by Wieland and Yang [2020]).[28] For each of

[27] 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 conference. Recent work highlights the role of the press conference in transmitting additional clarification and forward-looking guidance beyond the FOMC statement (see e.g., Gómez-Cram and Grotteria [2022], Swanson and Jayawickrema [2024], and Narain and Sangani [2026]).

[28] See Appendix Table A.1 for further details on each monetary policy surprise series. The Romer and Romer [2004]

these surprise measures, we create a quarterly-aggregated series by summing all surprises within a given quarter (as we did for our baseline statement surprise series).

Table 2 evaluates the issues of pre-trends (columns 1 and 2) and relevance (column 3) across these off-the-shelf measures of monetary policy surprises. To examine pre-trends, we estimate regressions of the form:

$$\Delta y^{(1)}_{t-k} = \alpha_k + \beta_k \cdot shock_t + \varepsilon_{t,k}, \quad k \in \{1, 2, \ldots, 8\}, \tag{2}$$

where $\Delta y^{(1)}_{t-k}$ is the change in the one-year Treasury yield from quarter $t-k-1$ to quarter $t-k$ and $shock_t$ is a quarterly-aggregated monetary policy surprise measure. In words, we examine whether surprises dated in quarter t are systematically related to yield movements realized *before* the surprise occurs. Thus, a non-zero β estimate indicates that the surprise series is correlated with prior movements in the one-year yield. Column 1 of Table 2 reports the p-value for the test $\beta_1 = 0$—corresponding to the most recent quarter prior to the surprise—while column 2 reports the p-value for the joint test $\beta_1 = \beta_2 = \ldots = \beta_8 = 0$. The low p-values reported across columns 1 and 2 indicate that low-frequency pre-trends are common across a range of monetary policy surprise measures (even orthogonalized measures).

Turning to relevance, column 3 of Table 2 reports the F-statistic for the specification in Equation (2) when $k = 0$, that is:

$$\Delta y^{(1)}_{t} = \alpha_0 + \beta_0 \cdot shock_t + \varepsilon_{t,0},$$

where $\Delta y^{(1)}_{t}$ is the change in the one-year Treasury yield from quarter $t-1$ to quarter t. In words, this regression measures the strength of the relationship between each quarterly-aggregated surprise series and the *contemporaneous* quarterly movement in the one-year yield.

Looking across columns 1, 2, and 3 clarifies that pre-trends and/or relevance issues exist broadly. In fact, *all* of the surprise measures examined in Table 2 exhibit significant pre-trends, or a low F-statistic, or both.

Table 3 takes this analysis one step further by checking if the addition of our own controls can ameliorate pre-trends while maintaining first-stage relevance. Similar to Section 2.2, we consider three sets of controls: (i) the set of pre-announcement macroeconomic and financial variables used by Bauer and Swanson [2023], (ii) four quarterly lags of real GDP growth, unemployment, and inflation, and (iii) the past eight lags of quarterly yield changes.[29]

To examine whether the first two sets of controls resolve the issue of pre-trends, we again test

series is a narrative measure rather than a high-frequency surprise; we include it here for completeness. We have also found similar results when experimenting with other shock series identified with the narrative approach (e.g., Aruoba and Drechsel [2026]) over the same sample period.

[29]We do not include the orthogonalized shocks from Acosta et al. [2025] and Bauer and Swanson [2023] in Table 3, since the column "Controlling for B&S controls" replicates both papers' orthogonalized shock series.

Table 2: Off-the-Shelf Monetary Policy Surprise Series: Examining Pre-Trends and Relevance

	Pre-trends		Relevance
	Coef. on last quarter yield change $\Delta y^{(1)}_{t-1}$	Coefs. on last eight quarters yield changes $\{\Delta y^{(1)}_{t-k}\}^{8}_{k=1}$	Effect on shock quarter yield change $\Delta y^{(1)}_{t}$
Shock	p-value (1)	Joint test p-value (2)	F-stat (3)
Raw shocks:			
Acosta et al. [2025] statement (*stmt*)	0.00	0.00	21.6
Acosta et al. [2025] monetary event (*me*)	0.00	0.00	13.0
Bauer and Swanson [2023] (*mps*)	0.00	0.00	35.4
Bu et al. [2021]	0.02	0.01	1.4
Gorodnichenko and Weber [2016]	0.00	0.01	7.8
Gürkaynak et al. [2005] target	0.05	0.04	0.0
Gürkaynak et al. [2005] path	0.00	0.00	31.4
Nakamura and Steinsson [2018]	0.00	0.00	15.6
Romer and Romer [2004]	0.00	0.00	6.0
Orthogonalized shocks:			
Acosta et al. [2025] statement ($stmt^{\perp}$)	0.00	0.02	0.3
Acosta et al. [2025] monetary event ($me^{\perp}$)	0.02	0.34	0.0
Bauer and Swanson [2023] ($mps^{\perp}$)	0.00	0.13	1.1
Miranda-Agrippino and Ricco [2021] (IV1)	0.04	0.20	4.1

Note: Column 1 reports the p-value that the coefficient $\beta_1 = 0$ using the specification in Equation (2) for $k = 1$. Column 2 reports the p-value of an analogous test that coefficients $\beta_1, \beta_2, \ldots, \beta_8$ are jointly zero. Column 3 reports the F-statistic for the specification in Equation (2) for $k = 0$. See Appendix Table A.1 for details on each monetary policy surprise series.

for pre-trends using specification (2), but now with controls. Columns 1 and 3 of Table 3 report, respectively, the p-value for the test that the coefficients on the past eight lags of yield changes are jointly zero, i.e., $\beta_1 = \beta_2 = \ldots = \beta_8 = 0$, in the specification augmented with each set of controls. To assess the relevance of the shock after controls are added, columns 2 and 4 of Table 3 report the F-statistic for the specification in Equation (2) when $k = 0$ and with the corresponding controls included. Echoing the results in Table 2, we again find—regardless of whether we use the Bauer and Swanson [2023] controls or the lags of macroeconomic variables as controls—that *all* of the monetary policy surprise series in Table 3 exhibit significant pre-trends, a low F-statistic, or both.

Including eight lags of yield changes as controls (the third set of controls) mechanically ensures that $\beta_1 = \beta_2 = \ldots = \beta_8 = 0$. To check for pre-trends in this case, we instead regress each of the Bauer and Swanson [2023] controls from quarter $t-1$ on the time-t monetary policy surprise, including the lagged yield changes as controls. Column 5 of Table 3 reports the p-value for the test that the coefficient on the surprise series is jointly zero across controls. In column 6 of Table 3, our calculation of the F-statistic mirrors the approach in columns 2 and 4. Again, for each monetary policy surprise series, we find either significant pre-trends (column 5) or a low F-statistic (column 6) after adding eight lags of yield changes as controls.

Overall, Tables 2 and 3 illustrate across a wide range of monetary policy surprise measures that quarterly-aggregated surprises are either systematically related to prior movements in the one-year Treasury yield, weakly related to contemporaneous yield movements, or both. Thus, the central concern is not simply that a particular high-frequency surprise measure is contaminated, but rather that any empirical mapping from high-frequency monetary policy surprises into lower-frequency macroeconomic outcomes is likely to introduce business-cycle variation that is difficult to purge without also removing much of the policy-relevant variation.

2.4 Using High-Frequency Monetary Policy Surprises at High Frequency

As discussed in Section 2.1, the business-cycle endogeneity problem creates an identification concern for empirical designs that use monetary policy surprises to estimate low-frequency dynamics. By contrast, high-frequency event studies—which focus on high-frequency outcomes—are likely to be less exposed to the low-frequency pre-trends that we uncover. This is because high-frequency identification rests on the comparatively narrow assumption that monetary policy surprises are the primary driver of outcome movements only within tight event windows. Of course, high-frequency event studies rely on the outcome of interest actually being observable in a tight window around FOMC announcements. While this is not the case for many real macroeconomic variables (e.g., investment), it is typically true for asset prices.

To assess the validity of high-frequency event studies, Figure 5 presents daily event-study estimates of the effect of monetary policy statement surprises on four different asset prices: the one-year yield, the S&P 500, the 10-year yield, and the 10-year real yield. In each panel, we present the raw estimate as well as the corresponding estimate after including each of the four sets of low-frequency controls used in Figures 3 and 4. In stark contrast to the sensitivity at quarterly frequencies that was illustrated in Figures 3 and 4, Figure 5 instead shows that estimated asset-price responses are insensitive to the inclusion of controls in narrow windows around a monetary policy surprise. For example, in a two-day window spanning from the close of the trading day preceding the FOMC meeting ($t-1$) to the close of the trading day following the meeting ($t+1$), estimated asset-price responses are almost completely insensitive to the controls used (we focus on this two-day window for consistency with the estimation of our Q-theory-based framework in Section 4; further details therein).

To make this point more formally, Table 4 conducts an Oster [2019] exercise to quantify the sensitivity of estimated asset-price responses in a narrow two-day window around monetary pol-

Table 3: Monetary Policy Surprise Series with Controls: Examining Pre-Trends and Relevance

	Controlling for B&S controls		Controlling for 4 lags of macro vars		Controlling for 8 lags of yield changes	
	Pre-trends	Relevance	Pre-trends	Relevance	Pre-trends	Relevance
	Coefs. on $\{\Delta y_{t-k}^{(1)}\}_{k=1}^{8}$	Effect on $\Delta y_t^{(1)}$	Coefs. on $\{\Delta y_{t-k}^{(1)}\}_{k=1}^{8}$	Effect on $\Delta y_t^{(1)}$	Coefs. on B&S cntrls	Effect on $\Delta y_t^{(1)}$
Shock	p-value (1)	F-stat (2)	p-value (3)	F-stat (4)	p-value (5)	F-stat (6)
Acosta et al. [2025] statement ($stmt$)	0.02	0.1	0.00	14.8	0.00	6.8
Acosta et al. [2025] monetary event (me)	0.28	0.3	0.00	10.1	0.00	3.5
Bauer and Swanson [2023] (mps)	0.07	3.6	0.00	25.1	0.00	19.8
Bu et al. [2021]	0.04	0.3	0.00	2.2	0.28	0.5
Gorodnichenko and Weber [2016]	0.29	0.5	0.11	3.7	0.01	4.7
Gürkaynak et al. [2005] target	0.23	3.8	0.04	0.0	0.49	1.2
Gürkaynak et al. [2005] path	0.00	7.0	0.00	22.9	0.00	25.2
Miranda-Agrippino and Ricco [2021] (IV1)	0.46	0.6	0.17	4.3	0.35	1.6
Nakamura and Steinsson [2018]	0.10	0.1	0.00	10.2	0.00	6.2
Romer and Romer [2004]	0.00	0.0	0.00	3.0	0.27	0.9

Note: Columns 1–2 report results for each surprise series when including the controls used by Bauer and Swanson [2023]. Columns 3–4 report results for each surprise series when controlling for four lags of real GDP growth, inflation, and unemployment. Columns 5–6 report results for each surprise series when controlling for eight lags of quarterly yield changes, $\{\Delta y_{t-k}^{(1)}\}_{k=1}^{8}$. See the text for further description. See Appendix Table A.1 for details on each monetary policy surprise series.

Figure 5: Path of Asset Prices Around Monetary Policy Statement Surprise: Daily Frequency

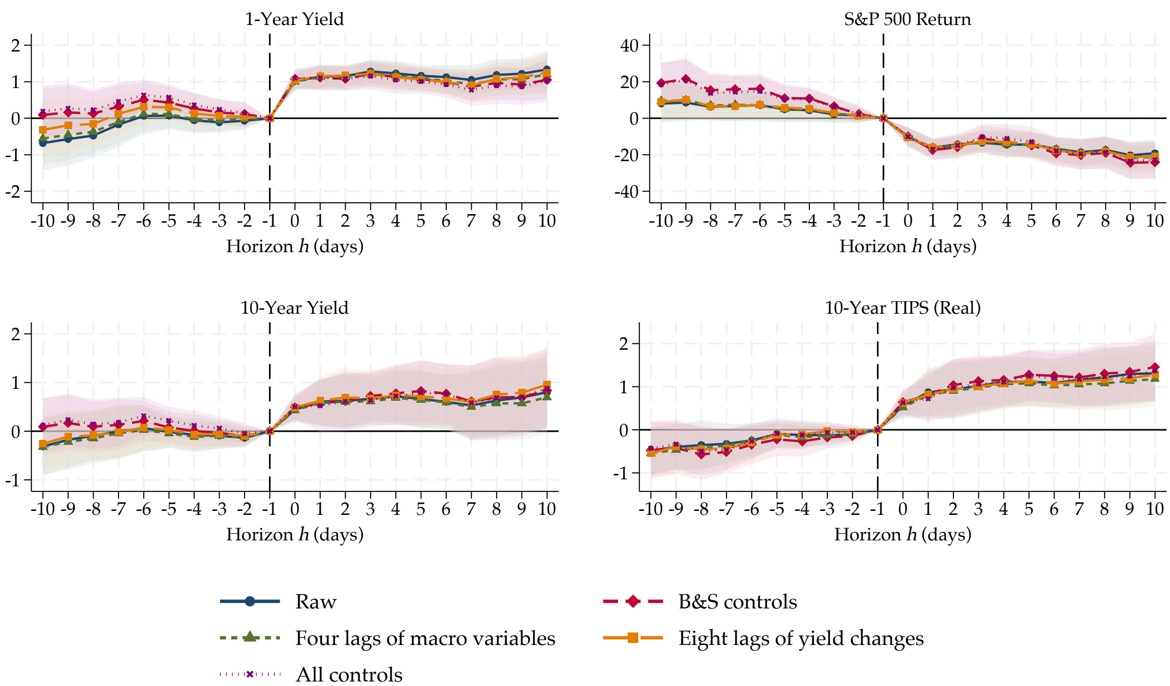

Note: This figure plots the sequence $\{\hat{\beta}_h^v\}_{h=-10}^{10}$ estimated from the specification

$$x_{t+h} - x_{t-1} = \alpha_h^v + \beta_h^v \cdot stmt_t + \gamma_h^v X_t^v + \varepsilon_{t,h}^v,$$

where x_t is an outcome x on day t, $stmt_t$ is the statement surprise from Acosta et al. [2025] on day t, and X^v is a set of controls. For "S&P 500 Return," the outcome is the cumulative S&P 500 return, normalized to zero on day $t-1$. The yield and S&P 500 panels use the full sample of 242 statement surprises over 1994–2023; the 10-year TIPS panel begins in 1999, when the real-yield curve becomes available, and uses 200 surprises. Each line corresponds to a different set of variables v. B&S controls are the controls used by Bauer and Swanson [2023]: the most recent nonfarm payrolls surprise, employment growth, stock-market returns, yield-curve slope changes, commodity-price changes, and option-implied Treasury skewness. Macro controls are four lags of real GDP growth, inflation, and unemployment. Lags of yield changes are quarterly changes in the one-year yield in the eight quarters preceding the surprise. Confidence intervals are constructed at the 90% level with heteroskedasticity-robust standard errors.

icy surprises. Again in contrast to the lower-frequency exercise in Table 1, Table 4 now shows tight bounds on our high-frequency (two-day) estimates of the effect of monetary policy surprises on asset prices. Overall, the analysis in this subsection suggests that high-frequency event studies are robust to the lower-frequency pre-trends that otherwise contaminate estimates based on monthly or quarterly outcomes.

2.5 Summary and Implications for Future Research

The evidence in this section points to a fundamental limitation of empirical designs that use high-frequency monetary policy surprises to identify the response of lower-frequency out-

Table 4: Sensitivity of Estimated Effects on Daily Asset Prices: Oster [2019] Bounds

	Effect of monetary surprise on outcome from $t = -1$ to $t = +1$ (pp)					
			Controls for 8 lags of yield changes		All controls	
Outcome	$\hat{\beta}_{\text{short}}$	R^2_{short}	R^2_{long}	Oster [2019] bounds	R^2_{long}	Oster [2019] bounds
1-year Treasury yield	1.14	0.296	0.340	[1.13, 1.18]	0.390	[1.07, 1.15]
5-year Treasury yield	0.89	0.095	0.141	[0.90, 1.04]	0.196	[0.91, 0.95]
10-year Treasury yield	0.61	0.042	0.075	[0.61, 0.64]	0.119	[0.49, 0.56]
5-year TIPS yield	1.20	0.118	0.174	[1.06, 1.20]	0.265	[0.97, 1.13]
10-year TIPS yield	0.87	0.085	0.149	[0.77, 0.85]	0.220	[0.65, 0.80]
S&P 500 return	-16.08	0.114	0.120	[-16.15, -15.97]	0.195	[-17.57, -16.32]

Note: The columns $\hat{\beta}_{\text{short}}$ and R^2_{short} report results from $x_{t+1} - x_{t-1} = \beta_{\text{short}} \cdot stmt_t + \varepsilon_t$ for each outcome x. For "S&P 500 Return," the outcome is the cumulative S&P 500 return, normalized to zero on day $t-1$. The yield and S&P 500 rows use 242 statement surprises over 1994–2023; the TIPS rows begin in 1999, when the real-yield curve becomes available, and use 200 surprises. The two pairs of long columns differ in their control set: the first ("8 lags of yield changes") adds only the eight quarterly yield-change lags $\{\Delta y^{(1)}_{t-k}\}_{k=1}^{8}$; the second ("All controls") adds those plus four quarterly lags of GDP, unemployment, and inflation, and the B&S controls. The Oster [2019] range is calculated by varying $\delta \in [-1, 1]$ using $\beta^* = \beta_{\text{long}} - \delta \left(\beta_{\text{short}} - \beta_{\text{long}}\right) \frac{R^2_{\text{max}} - R^2_{\text{long}}}{R^2_{\text{long}} - R^2_{\text{short}}}$, where R^2_{long} is the R-squared of the regression with controls and $R^2_{\text{max}} = \min\{1, 1.3 \cdot R^2_{\text{long}}\}$.

comes. Across a wide range of commonly used surprise measures, we document that quarterly-aggregated surprises are systematically related to prior yield movements (or weakly related to contemporaneous yield movements). This pattern suggests that estimated responses of lower-frequency outcomes may partly reflect the continuation of pre-existing business-cycle dynamics rather than the causal effect of monetary policy alone.

Our findings have two implications for empirical research that maps high-frequency monetary policy surprises into lower-frequency outcomes. The first is diagnostic: empirical work using monthly- or quarterly-aggregated surprises should examine not only first-stage relevance, but also whether the proposed surprise series is systematically related to pre-existing business-cycle conditions. Researchers can further evaluate the scope of these concerns using the sensitivity and omitted-variable-bias analyses applied above. The second is methodological: because high-frequency event studies appear robust to these business-cycle endogeneity concerns, empirical designs that preserve the high-frequency identifying variation in monetary policy surprises while avoiding the contamination introduced by directly mapping those surprises into lower-frequency outcomes may offer a more credible path forward. In the remainder of the paper, we pursue one such approach in the context of corporate investment.

3 Monetary Policy and Corporate Investment: Sufficient Statistics in a Q-Theory-Based Framework

The business-cycle endogeneity problem that we detailed in Section 2 suggests the need for alternative empirical designs to evaluate the impacts of monetary policy. Although we are unable to provide a general solution to this business-cycle endogeneity problem that applies for all outcome variables, we do propose a path forward in the context of corporate investment.[30] This Q-theory-based methodology for estimating the corporate investment response to monetary policy will be our focus in the remainder of this paper.

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. Moreover, we will show in Section 4.4 that Q-theory also has strong explanatory power for firm-level investment in our sample of U.S. publicly traded firms.[31] In particular, we find that our estimated firm-level investment-Q sensitivities predict firm investment out-of-sample with an R^2 of roughly 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 its 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 matter for aggregation, and which future research should continue to explore.

[30] As summarized in footnote 11, the empirical design of regressing investment on quarterly-aggregated monetary policy surprises is particularly common in the corporate-investment context, suggesting the usefulness of alternative empirical approaches such as the one we propose here.

[31] As mentioned in footnote 7, we make no claims about Q-theory's explanatory power for private firms.

3.1 Q-Theory-Based Framework

The framework proposed below is, for now, more of a sketch than a fully specified model. This will be refined in future drafts.

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.[32] 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:

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

In this framework, 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:[33]

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

To estimate how V_{it} responds to a monetary policy shock (dr_t) 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 shock:*

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

[32] This notation implicitly assumes that these time-t state variables are sufficient for characterizing the firm's value function V, including after a monetary policy shock of dr_t (as studied below). Although we will denote a monetary policy shock in this section by its contemporaneous effect on the interest rate, dr_t, this shock may also affect the expected future path of interest rates.

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

The assumption that $\frac{d\text{Loans}_{it}}{dr_t} \approx 0$ is justified on the grounds that most bank loans have floating interest rates, as emphasized for example by Ippolito et al. [2018].[34] While monetary policy shocks may still affect the price of credit risk and hence alter the market value of bank loans through general-equilibrium channels, Assumption 1 abstracts from such effects. The assumption that $\frac{d\text{Net Working Capital}_{it}}{dr_t} \approx 0$ is justified on the grounds that current assets and liabilities tend to be short-duration and/or floating-rate claims.[35]

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 shock dr_t simply by looking at the change in the value of firm i's W portfolio (which is observable using market-price data):

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

Aggregating over all firms, we can calculate the overall effect of a monetary policy shock on contemporaneous aggregate investment as follows:[36]

$$\frac{dI_t}{dr_t} = \sum_i \frac{dI_{it}}{dr_t} = \sum_i \gamma_i \frac{dV_{it}}{dr_t} = \sum_i \gamma_i \frac{dW_{it}}{dr_t} = \sum_i W_{it}\gamma_i \frac{dW_{it}/dr_t}{W_{it}},$$

where we multiply and divide by W_{it} in the last equality to express the change in firm-level W portfolios as a percentage response, $\frac{dW_{it}/dr_t}{W_{it}}$. We will refer to firm i's response $\frac{dW_{it}/dr_t}{W_{it}}$ as its *return sensitivity* to a monetary policy shock and, assuming it is time-invariant, denote it by β_i. We will estimate β_i as the firm-specific slope of W-portfolio returns on monetary policy surprises in Section 4.

Finally, denote by W_t the *aggregate* market value of all firms' outstanding equity and bonds at time t. Then, scaling the aggregate investment response by the aggregate W portfolio implies:

$$\frac{dI_t/dr_t}{W_t} = \sum_i \left(\frac{W_{it}}{W_t}\right)\gamma_i\beta_i = \mathbb{E}_w[\gamma_i] \times \mathbb{E}_w[\beta_i] + \mathbb{C}\text{ov}_w(\gamma_i, \beta_i), \tag{3}$$

where $\mathbb{E}_w[\gamma_i]$ is the value-weighted average investment-Q sensitivity (i.e., weighting by the rel-

[34] 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.

[35] 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.

[36] 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.

ative market value of each firm's W portfolio), $\mathbb{E}_w[\beta_i]$ is the value-weighted average return sensitivity to a monetary policy shock, and $\mathbb{C}\mathrm{ov}_w(\gamma_i, \beta_i)$ is the value-weighted covariance between individual firms' investment-Q sensitivity and their return sensitivity to a monetary policy shock. To avoid confusion, we quickly note that we generally expect Equation (3) to be negative—that is, a monetary policy tightening ($dr_t > 0$) will lead to a decrease in investment.

Equation (3) breaks down the aggregate corporate investment response 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, and (ii) the value-weighted average return sensitivity to a monetary policy shock. Intuitively, $\mathbb{E}_w[\beta_i]$ captures the sensitivity of the average firm's enterprise value to a monetary policy shock, while $\mathbb{E}_w[\gamma_i]$ captures the average firm's passthrough from a change in enterprise value to a change in investment.

The second term captures the role of firm-level heterogeneity. As Equation (3) highlights, the importance of firm-level heterogeneity depends on the covariance between individual firms' investment-Q sensitivity and the return sensitivity of firms' W portfolios to a monetary policy shock.[37] Intuitively, 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 higher investment-Q sensitivities.

3.2 Estimation Strategy

Equation (3) provides a Q-theory-based framework for estimating the corporate investment response to monetary policy surprises. Taking this framework to the data requires two separate steps. First, we need to estimate the return sensitivity of firm-level W portfolios around monetary policy surprises (i.e., β_i). 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 firm-level investment rates. 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 (3). We implement each of these steps in turn in Section 4.

[37] The role of covariances in micro-to-macro aggregation is well-known, and has been highlighted in recent work such as Auclert [2019] and Patterson [2023].

4 Taking Q-Theory-Based Framework to the Data

4.1 Data Description

Return Sensitivity to Monetary Policy Surprises. When estimating the return sensitivity of firm-level W portfolios to 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].

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 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 A 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 1994 to 2023 (excluding 2020).[38] 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.

[38] Although we observe corporate bond returns starting in 2002, we can make use of a longer time series when estimating firms' investment-Q sensitivities. We start in 1994 for consistency with when our monetary policy statement surprises become available, which is also the start of the equity-only, longer-sample robustness exercise reported in Section 4.6.

Analysis Sample. We require firms to have non-missing returns for at least 50 FOMC events to estimate return sensitivities to monetary policy surprises, and we require firms to have at least 10 years of investment and Q data to estimate investment-Q sensitivities. Our baseline analysis sample covers the approximately 2,500 firms for which we can estimate both sensitivities. As reported in Appendix Table A.2, this sample represents 86.6% of aggregate equity and bond value among all non-financial, non-utility Compustat firms. We also note that firms' W portfolios tend to be equity-dominant, with equity averaging 88% of portfolio value.

Robustness Checks: Deferred to Section 4.6. We now proceed to estimate the two sets of inputs, β_i and γ_i, needed to quantify the corporate investment response to monetary policy implied by Equation (3). Along the way, we will flag key robustness checks as needed, but will defer a fuller discussion until after presenting our main estimate of Equation (3) in Section 4.6 (with complete robustness results reported in Appendix Table A.4). By postponing a systematic robustness analysis until that stage, the reader will be better able to assess how various alternative choices affect our final aggregate-investment estimate.

4.2 Estimating Return Sensitivity to Monetary Policy Surprises

To bring the framework from Section 3 to the data, our first step is to leverage high-frequency identification to estimate the return on firm-level W portfolios around monetary policy surprises. We focus on the period 2002–2023, reflecting the availability of corporate bond data required to measure changes in enterprise value.

Firm-Level Return Sensitivity to Monetary Policy Surprises. To estimate firm-level return sensitivities, for each firm i in our sample we regress its event returns on the monetary policy surprise across FOMC events:

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

where R_{it} is the two-day cumulative return on firm i's W portfolio of stocks and bonds at event t, measured from the close of the trading day before the FOMC announcement (day -1) to the close of the day after (day $+1$), and $stmt_t$ is a monetary policy statement surprise [Acosta et al., 2025]; recall that $stmt_t$ is scaled so that one unit corresponds to a one percentage point high-frequency surprise in the one-year Treasury yield.

To construct the equity-and-bond return measure R_{it} in Equation (4), we measure the two-day cumulative return on equity and bonds separately, and then take a market-value-weighted average of the two returns (weighting each leg by its market value at the close of the trading day before the meeting). The equity return is the two-day CRSP return, inclusive of dividends;

the bond return is the two-day return across the firm's traded bonds, computed from TRACE transaction prices and inclusive of accrued interest and coupon.[39] Since corporate bonds can trade infrequently, we use the value-weighted return of a firm's traded bonds when at least one bond trades.[40] If a firm has no traded bonds, we impute the bond return using the value-weighted "market" bond return across the sample firms that do trade.[41]

We use a two-day window when measuring returns because cumulative price adjustments following a monetary policy surprise appear to not be fully realized on the announcement day alone. This pattern can clearly be seen in the S&P 500 panel of Figure 5. A comparable pattern also arises for the aggregate W portfolio, as illustrated in Appendix Figure A.5.[42,43] We will utilize a one-day window for robustness (the announcement-day return, from the close of day -1 to the close of day 0), and find that estimated return sensitivities are directionally consistent but generally smaller in magnitude.

Figure 6 illustrates our estimates of firm-level return sensitivities by showing a value-weighted histogram of the $\widehat{\beta}_i$ coefficients from Equation (4). The dashed red vertical line marks the (value-weighted) average return sensitivity to a one-unit monetary policy surprise; i.e., our estimate of $\mathbb{E}_w\left[\beta_i\right]$ in Equation (3). This average return sensitivity is -0.16—meaning that a one-unit surprise tightening lowers average-firm returns by 16%—which is similar in magnitude to the S&P 500 panel of Figure 5. Further summary statistics for our estimated firm-level return sensitivities, $\widehat{\beta}_i$, are provided in Appendix Table A.2.

While there will undoubtedly be some estimation noise in our firm-level $\widehat{\beta}_i$ estimates, a cross-

[39] Two-day returns are compounded from the two daily returns, which we winsorize separately for stocks and bonds at the median $\pm\, 5\times$IQR before compounding. We apply this same median $\pm\, 5\times$IQR winsorization to the estimated firm-level coefficients as well. Tightening the winsorization of the estimated firm-level coefficients to $3\times$IQR leaves the implied aggregate investment response essentially unchanged (see Appendix Table A.4).

[40] We count a bond as traded if it prints on day -1, day 0, and day 1 (i.e., the three days from which the two-day return is built).

[41] Note that equity plus actually-traded bonds cover about 95% of firm value (Appendix Table A.2), and the market bond-return imputation specifically applies to only about 1.2% of value. The implied aggregate investment response is correspondingly robust to the treatment of untraded bonds (see Appendix Table A.4).

[42] Recall that the aggregate W portfolio is the value-weighted portfolio of outstanding bonds and equities issued by all firms in our sample. The response of the aggregate W portfolio in Appendix Figure A.5 differs slightly from the value-weighted average firm-level return sensitivity $\mathbb{E}_w[\widehat{\beta}_i]$ that will soon be reported in Figure 6: the former is a single slope from a portfolio-level regression of the value-weighted W-portfolio return on the surprise, while the latter is a value-weighted average of individual firms' slopes. The two need not coincide because the aggregate W portfolio's weights move across events, and because the panel is unbalanced.

[43] For more, Appendix Figure A.6 shows the response of the value-weighted portfolio of equities alone. The resulting pattern is nearly identical to Appendix Figure A.5, reflecting the fact that, in our sample, equities account for 88% of the market value of the aggregate W portfolio. Appendix Figure A.7 reproduces the same event-study analysis for the value-weighted portfolio of corporate bonds in our sample. 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. Finally, Appendix Figure A.8 replicates the event study in Appendix Figure A.5 but with the addition of controls. Consistent with the analysis in Section 2.4, we again find that our high-frequency (two-day) return estimates exhibit minimal sensitivity to the inclusion of various controls.

Figure 6: Value-Weighted Histogram of Firm-Level Return Sensitivity Estimates

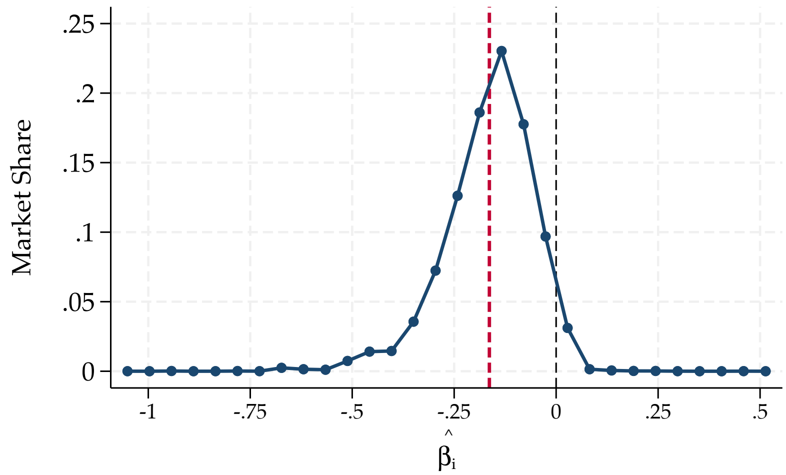

Note: This figure plots the value-weighted distribution of the firm-level return sensitivities $\widehat{\beta}_i$, estimated one slope per firm from Equation (4) over our 2002–2023 sample. Each firm is weighted by its average share of the sample's aggregate market value (equity plus bonds) across events, so the height of each bin equals the combined market-value share of the firms whose $\widehat{\beta}_i$ falls in it, and the bin heights sum to one. The dashed red line marks the value-weighted mean, $\mathbb{E}_w[\widehat{\beta}_i] = -0.16$; the thin dashed line marks zero. Return sensitivities are winsorized at the median $\pm\, 5\times$IQR.

validation exercise shows that they nonetheless capture meaningful heterogeneity in return sensitivity. We sort firms into bins by their $\widehat{\beta}_i$ estimated on a training subset of FOMC events, and then estimate the realized return-shock slope of each bin on held-out events that were *not* used to form the bins. As shown in Appendix Figure A.9 (which describes the procedure in full), firms in bins with more negative $\widehat{\beta}$ estimates also exhibit stronger sensitivity in the held-out events, substantiating the conclusion that our estimates capture persistent heterogeneity rather than noise.

We again note that robustness exercises are deferred until Section 4.6. For our firm-level return sensitivity estimates, we explore robustness to (among others): the addition of controls (echoing the analysis in Section 2.4), reweighting to reduce the influence of large firms, using the broader monetary policy event surprise measure (me_t), and using a one-day return window instead of our baseline two-day window. Using a one-day window roughly halves the investment response implied by Equation (3)—this attenuation follows from the S&P 500 panel of Figure 5, which shows that only about half of the overall return response occurs on the announcement day. Otherwise, these alternatives have only modest effects on our quantification of Equation (3).

4.3 Estimating Investment-Q Sensitivity

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

$$\frac{I_{it}}{K_{it-1}} = \alpha_i + \gamma_i \cdot Q_{it-1} + \delta_{jt} + \epsilon_{it}, \tag{5}$$

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).[44] A key motivation for these fixed effects is that measured (gross) investment reflects not only the firm's response to investment opportunities—captured by Q—but also the replacement of depreciated capital. That is, Q is a sufficient statistic for the investment rate only up to a depreciation (replacement) term. Because depreciation rates vary across firms and, through asset composition, across industries over time—and may be correlated with Q—firm fixed effects absorb persistent cross-firm differences in depreciation while industry×year fixed effects absorb sector-level movements in depreciation over time.

We evaluate two different versions of the investment ratio in Equation (5). The first uses a tangible investment ratio (capital expenditures scaled by lagged tangible assets).[45] The second uses a total investment ratio that includes intangibles, following Eisfeldt and Papanikolaou [2013] and Peters and Taylor [2017].[46]

To illustrate firm-level investment-Q sensitivities, Figure 7 presents value-weighted histograms of the $\widehat{\gamma}_i$ coefficients from Equation (5). In each panel, the dashed red vertical line marks the (value-weighted) average investment-Q sensitivity; i.e., our estimate of $\mathbb{E}_w[\gamma_i]$ in Equation (3). For standard (tangible) investment and total investment, respectively, the average investment-Q sensitivity is 0.017 and 0.031. To interpret these magnitudes, a one-unit increase in a firm's Tobin's Q is associated with a rise in its annual investment rate of 1.7 percentage points for standard investment and 3.1 percentage points for total investment. Further summary statistics for our estimated investment-Q sensitivities, $\widehat{\gamma}_i$, are provided in Appendix Table A.3.

While there will again be noise in firm-level $\widehat{\gamma}_i$ estimates, a cross-validation exercise shows that they capture meaningful heterogeneity in investment-Q sensitivity. Using a leave-one-year-out design, we sort firms into bins by their $\widehat{\gamma}_i$ estimated on the other years, and then estimate the realized investment-Q slope of each bin on the held-out year that was *not* used to form the bins.

[44] 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.

[45] 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`×`csho`), plus the book value of debt (Compustat items `dltt` + `dlc`), minus the firm's current assets (Compustat item `act`).

[46] 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.

Figure 7: Value-Weighted Histogram of Firm-Level Investment-Q Sensitivity Estimates

(a) Standard Investment (b) Total Investment

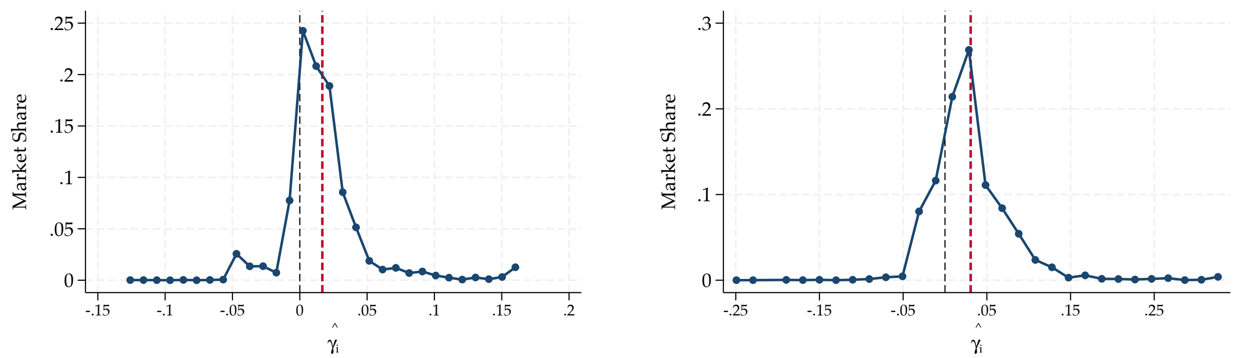

Note: This figure plots the value-weighted distribution of the firm-level investment-Q sensitivities $\widehat{\gamma}_i$ estimated from Equation (5); panel (a) uses the standard (tangible) investment measure and panel (b) the total-investment measure of Peters and Taylor [2017]. In each panel, firms are weighted by their average share of the sample's aggregate market value (equity plus bonds), so the height of each bin equals the combined market-value share of the firms whose $\widehat{\gamma}_i$ falls in it, and the bin heights sum to one. The dashed red line marks the value-weighted mean $\mathbb{E}_w[\widehat{\gamma}_i]$: 0.017 in panel (a) and 0.031 in panel (b); the thin dashed line marks zero. Sensitivities are winsorized at the median $\pm\,5\times$IQR.

As shown in Appendix Figure A.10 (which describes the procedure in full), firms in high-$\widehat{\gamma}$ bins exhibit stronger investment-Q sensitivity in the held-out year, substantiating the conclusion that our estimates capture persistent heterogeneity rather than noise.

When examining robustness in Section 4.6, we will explore (among others): reweighting to reduce the influence of large firms, re-estimating firm-level γ coefficients under coarser fixed-effects specifications, and estimating γ coefficients in industry$\times$size cells rather than at the firm level. These alternatives have only modest effects on our quantification of Equation (3).

4.4 Aside: Evaluating Firm-Level Q-Theory Performance

Because our sufficient-statistics framework relies on Q-theory as a model of firm investment, we now assess the explanatory power of Q-theory for firm-level investment in our sample.

Before presenting the results, we emphasize a distinctive feature of our implementation of Q-theory, which is that we estimate the γ_i coefficients separately for each firm. That is, we allow each firm to have its own investment-Q sensitivity. We refer to this empirical implementation as *firm-level Q-theory*. As we detail below, our firm-level implementation turns out to be quite important for improving Q-theory's explanatory power.

Starting with an in-sample evaluation of Q-theory for firm-level investment, Table 5 presents various R^2 metrics for the regression in Equation (5). We report each metric for both investment measures and for two specifications: our baseline heterogeneous-γ model (in which every firm has its own slope γ_i), and a homogeneous-γ benchmark that imposes a single common slope.

Both specifications include firm and industry×year fixed effects. In the heterogeneous-γ case the firm-specific slopes are estimated jointly in a single regression, and every R^2 we report is the adjusted $\bar{R}^2$ so that the additional slopes are penalized for the degrees of freedom they consume.

Table 5 decomposes the total R^2 into three related metrics. First, the restricted R^2 captures the baseline fit obtained from a fixed-effects-only regression that excludes the focal regression term, $\gamma_i \cdot Q_{it-1}$. Next, the incremental R^2 is defined as the difference between the total R^2 and the restricted R^2, thus capturing the explanatory power from adding the regression term $\gamma_i \cdot Q_{it-1}$. Finally, the partial R^2 rescales the incremental R^2 to capture the share of remaining variation accounted for by the added regression term. Specifically, the partial R^2 is defined as $\frac{\text{Incremental } R^2}{1-\text{Restricted } R^2}$, thereby rescaling the incremental R^2 by the residual variation left unexplained by the restricted model alone.

There are two main takeaways from Table 5. First, our firm-level Q-theory model accounts for a considerable portion of the variation in firm-level investment. Indeed, our baseline heterogeneous-γ specification achieves partial R^2 values of 29% and 40%, respectively, for standard and total investment. Second, by comparing the heterogeneous-γ columns to the homogeneous-γ columns, Table 5 illustrates the added explanatory power from allowing firms to have their own investment-Q sensitivity. The partial R^2 roughly doubles for both standard and total investment, implying that our firm-level implementation of Q-theory substantially improves its explanatory power.

Table 5: In-Sample Performance of Firm-Level Q-Theory

	Standard investment		Total investment	
	Het. γ	Hom. γ	Het. γ	Hom. γ
Total $\bar{R}^2$	53%	43%	69%	59%
Restricted $\bar{R}^2$	34%	34%	48%	48%
Incremental $\bar{R}^2$	19%	10%	21%	12%
Partial $\bar{R}^2$	29%	15%	40%	22%
F-test p-value	0.000		0.000	
Observations	51,697	51,697	51,697	51,697

Note: This table reports adjusted $\bar{R}^2$ statistics for the firm-level Q-theory investment regression in Equation (5), shown for standard (tangible) and total investment and, within each, for our baseline heterogeneous-γ specification (firm-specific slopes) and a homogeneous-γ benchmark (a single common slope). All regressions include firm and industry×year fixed effects. The restricted $\bar{R}^2$ comes from a fixed-effects-only model that omits the $\gamma_i\, Q_{it-1}$ term; the incremental $\bar{R}^2$ is the total minus the restricted; and the partial $\bar{R}^2$ rescales the incremental by the variation the restricted model leaves unexplained. The F-test p-value (heterogeneous-γ columns) is for the joint significance of the firm-specific slopes against the homogeneous benchmark. The regressions use firm-year observations from our baseline analysis sample, which is defined in Section 4.1.

Next, we look out of sample to evaluate the extent to which our firm-level $\widehat{\gamma}_i$ coefficients predict firm investment. To do so, we perform a different version of a leave-one-year-out exercise, as follows. For each year t in our sample, we estimate firm-specific slopes $\widehat{\gamma}_i$ from Equation (5) after excluding all year-t investment observations from the estimation. We also compute training-sample means $\bar{I}_i$ and $\bar{Q}_i$ for each firm using only the training years. We require at least 10 years of non-missing data per firm.

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, $\widehat{\gamma}_i \times (Q_{it-1} - \bar{Q}_i)$.[47] Specifically, we stack all held-out year observations and plot binned scatter plots of actual demeaned investment against predicted demeaned investment. Importantly, the $\widehat{\gamma}_i$ 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 8 displays the results. Panel (a) uses standard investment and 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 22% and 31% indicate substantial explanatory power. This justifies our use of Q-theory as a simple—but nonetheless realistic—model of firm investment.

For more, Appendix Figure A.11 provides the corresponding analysis under the restriction that firm-level γ coefficients are homogeneous. The out-of-sample fit clearly suffers when firm-level heterogeneity is removed.

4.5 Estimating Sensitivity Covariance

The next step in quantifying Equation (3) is to combine our estimates of firms' return sensitivities with firms' investment-Q sensitivities in order to estimate the covariance term, $\mathbb{C}\mathrm{ov}_w\left(\gamma_i, \beta_i\right)$, which captures the role of firm-level heterogeneity. In particular, Equation (3) implies that the investment response to monetary policy will be amplified if the firms with larger-magnitude return sensitivities (i.e., more negative) are also the firms with higher investment-Q sensitivities; that is, if $\mathbb{C}\mathrm{ov}_w\left(\gamma_i, \beta_i\right) < 0$.

We evaluate this covariance channel by estimating:

$$\widehat{\beta}_i = a + b \cdot \widehat{\gamma}_i + \epsilon_i, \tag{6}$$

where $\widehat{\beta}_i$ is firm i's return sensitivity to monetary policy surprises and $\widehat{\gamma}_i$ is its investment-Q sensitivity. All regressions are weighted by the firm's average share of the aggregate W portfolio across FOMC events. Because both $\widehat{\beta}_i$ and $\widehat{\gamma}_i$ are estimated rather than observed, we compute

[47] When predicting investment, we do not include industry$\times$year fixed effects because doing so would require using outcomes from the held-out year.

Figure 8: Out-of-Sample Investment Predictability

(a) Standard Investment (b) Total Investment

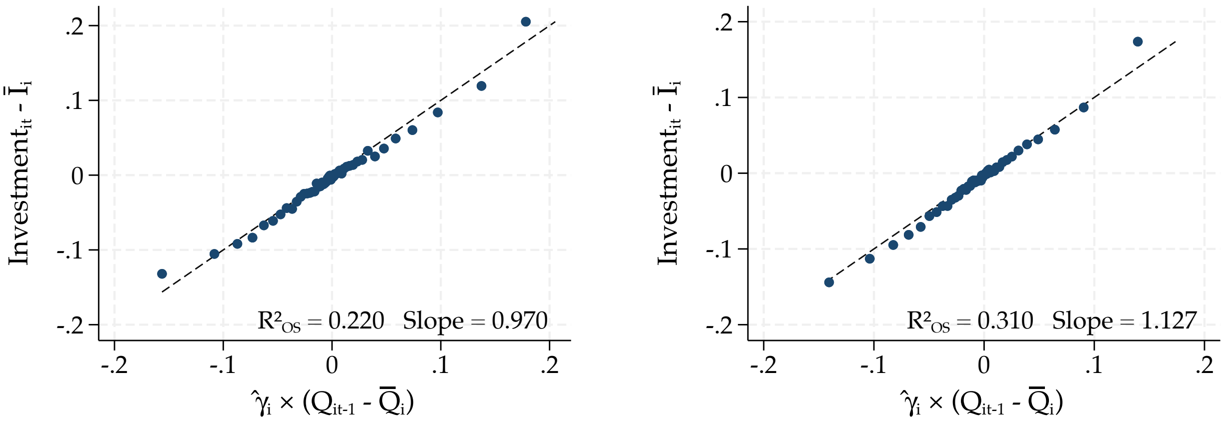

Note: Each panel shows the out-of-sample fit of firm-level Q-theory from the leave-one-year-out exercise described in the text. For each held-out year we predict a firm's demeaned investment as $\widehat{\gamma}_i \times (Q_{it-1} - \bar{Q}_i)$, where $\widehat{\gamma}_i$ is the firm's slope estimated on the remaining years and $\bar{I}_i$, $\bar{Q}_i$ are training-sample means; actual demeaned investment is $I_{it}/K_{it-1} - \bar{I}_i$. We winsorize the actual and the predicted demeaned investment at the median $\pm\, 5 \times$IQR, pool all held-out observations, sort them into 50 equal-sized bins by predicted investment, and plot bin averages of actual (y-axis) against predicted (x-axis) investment. The dashed line is the 45-degree line; the reported out-of-sample R^2 and slope are computed from the underlying, unbinned observations. Panel (a) uses standard (tangible) investment and panel (b) total investment.

standard errors via bootstrap.[48]

We also note that although $\widehat{\beta}_i$ and $\widehat{\gamma}_i$ are likely estimated with noise, such noise should not bias our estimate of the covariance term as long as the noise is not systematically related across the two estimates.[49] Since our firm-level coefficients $\widehat{\beta}_i$ and $\widehat{\gamma}_i$ are estimated from different data sources and rely on distinct identifying variation, this mitigates the concern that our covariance estimate primarily reflects correlated estimation error rather than true underlying heterogeneity.[50]

Before presenting regression estimates for Equation (6), Figure 9 provides a visualization by showing binned scatter plots of return sensitivity against investment-Q sensitivity. We sort firms into 50 bins based on their $\widehat{\gamma}_i$ estimate and plot the value-weighted average $\widehat{\beta}_i$ within each bin against the value-weighted average $\widehat{\gamma}_i$. 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. Both panels reveal a clear negative relationship: firms in

[48] We resample firms with replacement, re-estimate both $\widehat{\beta}_i$ and $\widehat{\gamma}_i$ in each bootstrap iteration, and run regression (6) on the bootstrapped sample. The standard deviation of the coefficient across bootstrap replications provides a standard error that accounts for the generated regressor problem.

[49] More fully, denote $\widehat{\beta}_i = \beta_i + \epsilon_i^\beta$ and $\widehat{\gamma}_i = \gamma_i + \epsilon_i^\gamma$. Then, our estimate $\mathbb{C}\text{ov}_w(\widehat{\gamma}_i, \widehat{\beta}_i)$ approximates the true covariance $\mathbb{C}\text{ov}_w(\gamma_i, \beta_i)$ as long as $\mathbb{C}\text{ov}_w(\gamma_i, \epsilon_i^\beta) \approx \mathbb{C}\text{ov}_w(\epsilon_i^\gamma, \beta_i) \approx \mathbb{C}\text{ov}_w(\epsilon_i^\gamma, \epsilon_i^\beta) \approx 0$.

[50] In addition, our covariance estimate generally remains stable across the robustness checks reported in Appendix Table A.4, including the specification that uses a bootstrap-adjusted covariance estimate.

higher-γ bins also tend to have returns that are more sensitive to monetary policy surprises.

Figure 9: Return Sensitivity and Investment-Q Sensitivity

(a) Standard Investment (b) Total Investment

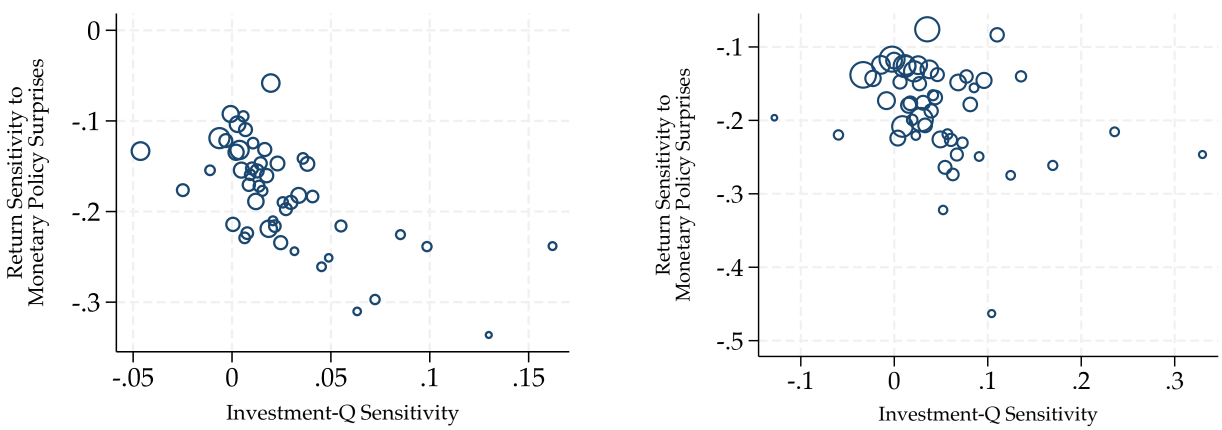

Note: Each panel plots value-weighted bin averages of return sensitivity to monetary policy ($\widehat{\beta}_i$, y-axis) against investment-Q sensitivity ($\widehat{\gamma}_i$, x-axis). Firms are sorted into 50 bins based on $\widehat{\gamma}_i$. Point size reflects the total W-portfolio value of the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

Now turning to the regression estimates for Equation (6), Table 6 presents our main result on the role of firm-level heterogeneity in shaping the corporate investment response to monetary policy. For both standard investment and total investment, the regression coefficient is negative and statistically significant, indicating that firms whose investment is more responsive to changes in Q also have a W portfolio that is more sensitive to monetary policy surprises. The bottom row of Table 6 reports the weighted covariance between $\widehat{\gamma}_i$ and $\widehat{\beta}_i$, computed as $\mathbb{C}\text{ov}_w(\widehat{\gamma}_i, \widehat{\beta}_i) = \hat{b} \times \text{Var}_w(\widehat{\gamma}_i)$. From Equation (3), this negative covariance implies an amplification effect stemming from firm-level heterogeneity.

Table 6: Return Sensitivity and Investment-Q Sensitivity: Covariance

	Standard investment	Total investment
$\hat{\gamma}_i$	-0.90***	-0.34***
	(-6.12)	(-3.08)
Constant	-0.15***	-0.15***
	(-18.63)	(-19.79)
Observations	2,517	2,517
R^2	0.064	0.021
$100 \times \mathbb{C}\text{ov}_w\left(\widehat{\gamma}_i, \widehat{\beta}_i\right)$	-0.09	-0.08

Note: This table reports estimates from cross-sectional regressions of return sensitivity ($\widehat{\beta}_i$) on investment-Q sensitivity ($\widehat{\gamma}_i$). 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: investment-Q sensitivity estimated using standard investment and standard Q. Column 2: investment-Q sensitivity estimated using total investment and total Q. $\mathbb{C}\text{ov}_w(\widehat{\gamma}_i, \widehat{\beta}_i)$ is the value-weighted covariance, computed as coefficient $\times$ $\text{Var}_w(\widehat{\gamma}_i)$.

4.6 Quantification: Corporate Investment Response to Monetary Policy

We are now ready to quantify the effect of monetary policy on corporate investment implied by Equation (3). Recall that:

$$\frac{dI_t/dr_t}{W_t} = \mathbb{E}_w\left[\gamma_i\right] \times \mathbb{E}_w\left[\beta_i\right] + \mathbb{C}\text{ov}_w\left(\gamma_i, \beta_i\right).$$

Using our baseline specification for standard investment, we estimate that $\mathbb{E}_w\left[\gamma_i\right] = 0.017$ (Figure 7), $\mathbb{E}_w\left[\beta_i\right] = -0.16$ (Figure 6), and $\mathbb{C}\text{ov}_w\left(\gamma_i, \beta_i\right) = -0.00085$ (Table 6), so that:

$$\frac{dI_t/dr_t}{W_t} = 0.017 \times -0.16 + (-0.00085) = -0.36\%.$$

Our analysis therefore implies that a one-unit monetary policy tightening reduces tangible investment by 0.36% of the value of the aggregate W portfolio (0.58% for total investment). Of this effect, about 24% (13% for total investment) is driven by the firm-level heterogeneity captured by the covariance term.

For interpretability, these investment responses can also be scaled relative to GDP and relative to firm capital. For GDP, since the aggregate W portfolio is roughly 0.9-times U.S. GDP over our sample period, we estimate that a one-unit tightening corresponds to an annual investment reduction of about 0.33% of U.S. GDP (0.53% for total investment).[51] Recall that this investment

[51] Although we find the scaling relative to GDP helpful for intuition, we nonetheless note that neither of our investment measures corresponds one-for-one with investment as defined in the national accounts.

response only corresponds to the firms in our sample, which account for 86.6% of the market value of all non-financial, non-utility Compustat firms. Applying the same per-dollar investment response to all non-financial, non-utility Compustat firms would imply an annual investment reduction of about 0.38% of U.S. GDP.

For capital, since the aggregate W portfolio is roughly 2.59-times the gross book value of physical capital for the firms in our sample, we estimate a fall in the aggregate investment *rate* (I/K) of 0.93% (0.76% for total investment).[52]

Finally, it is worth noting that our GDP and capital scalings are sensitive to the valuation date, as corporate market values have grown substantially over our sample period. Measured at end-of-sample (2023) valuations—when the aggregate W portfolio had risen to about 1.1-times U.S. GDP and 3.28-times the gross book value of physical capital—the same per-dollar investment response would instead correspond to 0.40% of U.S. GDP and 1.18% of capital.[53]

Robustness. In Appendix Table A.4 we repeat the quantification of Equation (3) under a broad range of alternative estimation assumptions, which we describe below. Each row of Appendix Table A.4 varies only the ingredient(s) named in the row or panel heading, and holds the rest at baseline. Across all of these alternatives, the implied aggregate response for tangible investment stays between roughly -0.14% and -0.50% of GDP. This suggests that our conclusions are not an artifact of any single modeling choice.

We start by examining robustness to our return-sensitivity estimates. In Panel A of Appendix Table A.4—which echoes the analysis in Section 2.4—we confirm that our high-frequency $\widehat{\beta}_i$ estimates are not sensitive to the addition of the controls used therein.

In Panel B, we examine robustness to our monetary surprise measure by replacing the statement-surprise series with the broader monetary-event surprise series (me_t) of Acosta et al. [2025], and by dropping all financial-crisis events (August 2007–June 2009). Next, we measure returns over a one-day window rather than our baseline two-day window. Finally, we estimate returns using an equity-only sample in order to exploit the statement-surprise series' availability back to 1994 (recall that our corporate bond data starts in 2002).[54]

In Panel C, we revisit how we handle the return on non-traded corporate bonds (see Section 4.2 for details on our baseline bond-return measure). First, rather than using the value-weighted market bond return for firms with no traded bonds, we now treat the bonds of no-trade firms as unchanged in value. Second, we impute nothing at all, and hold every untraded bond's value

[52] The aggregate W portfolio is roughly 1.32-times total capital for the intangibles-inclusive measure.

[53] This comparison presumes that the estimated sensitivities are stable over time, an assumption for which we provide a preliminary assessment in Appendix Table A.4 via an early- versus late-sample split—our estimates change only modestly across the two periods.

[54] Specifically, we assume that bond values are unchanged around announcements and we weight firms by their equity market value alone.

fixed.

With the exception of the one-day return window, our estimates across Panels A, B, and C exhibit minimal sensitivity. The one-day return window roughly halves the implied investment response, which—as discussed in Section 4.2—is natural given that only about half of the overall return response occurs on the announcement day.

Next, Panel D of Appendix Table A.4 studies the robustness of our investment-Q sensitivity estimates by perturbing the fixed-effects structure of the investment-Q regression in Equation (5). First, we replace the industry×year fixed effects with year fixed effects only. Second, we drop the year fixed effects as well.

Panel D suggests that the fixed-effects structure does have some impact on our investment estimate, albeit a modest one, with each of these coarser fixed-effects specifications yielding a somewhat stronger investment response.

Finally, Appendix Table A.4 reports robustness to assumptions that apply jointly to our return-sensitivity estimates and our investment-Q sensitivity estimates. In Panel E, we replace our firm-level estimates of β and γ with estimates computed at the Fama–French-49×size-decile cell level, which is a less granular unit of analysis.

In Panel F, we relax our baseline sample restrictions by lowering the floor on FOMC meetings (which binds on β) from 50 to 25, by lowering the floor on investment-Q years (which binds on γ) from 10 to 5, and then by estimating each sensitivity on its own largest sample by applying the meeting-count floor only to β and the investment-year floor only to γ, so that neither restriction constrains the other's estimate.[55]

Panel G gathers the remaining checks. The first four rows of Panel G focus on reducing the influence of noisily estimated sensitivities and heavily weighted firms. We start by tightening every winsorization from $5\times$IQR to $3\times$IQR. Next, we shrink each $\widehat{\beta}_i$ and $\widehat{\gamma}_i$ toward its cross-sectional mean using an empirical Bayes procedure, so that imprecisely estimated coefficients are pulled more strongly toward the mean.[56] Third, the value-weighted nature of our aggregates makes them potentially sensitive to a small number of very large firms (i.e., a granularity concern). We therefore cap each firm's value-share at the 95th percentile of the weight distribution (winsorizing the upper 5% of weights), renormalize so the weights sum to one, and re-estimate the aggregate moments. Fourth, we apply both adjustments simultaneously—shrinking the sensitivity estimates and winsorizing the firm weights—which jointly guards against the most problematic case of a heavily weighted firm whose sensitivity is also noisily estimated. Panel G then turns to

[55] In this case, $\mathbb{E}_w[\widehat{\beta}_i]$ and $\mathbb{E}_w[\widehat{\gamma}_i]$ are each computed on their own sample, with weights renormalized within that sample. The covariance term is computed for the firms that carry both estimates.

[56] Letting $\widehat{\beta}_i$ denote the OLS estimate and $\widehat{\sigma}_i$ its standard error, the shrunk estimate is $\lambda_i\widehat{\beta}_i + (1-\lambda_i)\bar{\beta}$, where $\lambda_i = \widehat{\tau}^2/(\widehat{\tau}^2 + \widehat{\sigma}_i^2)$ depends on the ratio of cross-sectional variance $\widehat{\tau}^2$ to estimation variance $\widehat{\sigma}_i^2$. The shrunk estimate for $\widehat{\gamma}_i$ is calculated similarly.

examining time variation by estimating the investment response on the early and late halves of the sample, respectively. Finally, we address the concern that the covariance term we estimate could suffer from a finite-sample bias if $\widehat{\beta}_i$ and $\widehat{\gamma}_i$ are estimated with error. We gauge this bias with a year-level cluster bootstrap that re-estimates both sensitivities on resampled years: the difference between the mean bootstrap covariance and the observed covariance estimates this bias, which we then subtract from the observed covariance. The bias-adjusted covariance leaves the aggregate response essentially unchanged, validating that the heterogeneity channel is not an artifact of estimation noise.

Across all of the joint robustness exercises in Panels E, F, and G, we again see that our investment estimates tend to vary only modestly, and the qualitative interpretation of our results appears robust across all robustness checks.

5 Conclusion

This paper makes two main contributions. First, we document a business-cycle endogeneity problem that arises when high-frequency monetary policy surprises are used outside of the narrow announcement windows in which they are constructed. Specifically, we find that these surprises are systematically related to pre-existing interest-rate dynamics, which undermines the causal interpretation of monetary policy event studies that use these surprises to estimate lower-frequency responses.

Second, we develop and implement a novel Q-theory-based framework for estimating the corporate investment response to monetary policy. This framework preserves high-frequency identification when estimating the effect of monetary policy on firm valuations, and then uses Q-theory to map from valuation changes to investment changes. Our baseline estimate implies that a one-unit monetary policy tightening reduces annual (tangible) investment by 0.33% of GDP for the firms in our sample, or roughly 0.38% of GDP if extrapolated to all non-financial, non-utility public firms. In future research we intend to investigate why the firms with larger return sensitivities 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

Boaz Abramson, Pablo De Llanos, and Lu Han. Monetary policy and rents. *Mimeo*, 2026.

Miguel Acosta, Connor M. Brennan, and Margaret M. Jacobson. Constructing high-frequency monetary policy surprises from SOFR futures. *Economics Letters*, 242:111873, 2024.

Miguel Acosta, Andrea Ajello, Michael D. Bauer, Francesca Loria, and Silvia Miranda-Agrippino.

Financial market effects of FOMC communication: Evidence from a new event-study database. *Federal Reserve Bank of San Francisco Working Paper 2025-30*, 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.

Boragan Aruoba and Thomas Drechsel. Identifying monetary policy shocks: A natural language approach. *American Economic Journal: Macroeconomics*, 2026. Forthcoming.

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

S. Mahdi Barakchian and Christopher Crowe. Monetary policy matters: Evidence from new shocks data. *Journal of Monetary Economics*, 60(8):950–966, 2013.

Michael D. Bauer and Mikhail Chernov. Interest rate skewness and biased beliefs. *Journal of Finance*, 79(1):173–217, 2024.

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.

Michael D. Bauer, Carolin E. Pflueger, and Adi Sunderam. Perceptions about monetary policy. *Quarterly Journal of Economics*, 139(4):2227–2278, 2024.

Mehdi Beyhaghi, Murray Z. Frank, Ping McLemore, and Ali Sanati. How do public and private firms respond to monetary policy shocks? *SSRN 4793393*, 2026.

Florin O. Bilbiie and Diego R. Känzig. Greed? Profits, inflation, and aggregate demand. *NBER Working Paper 31618*, 2023.

Chunya Bu, John Rogers, and Wenbin Wu. A unified measure of Fed monetary policy shocks. *Journal of Monetary Economics*, 118:331–349, 2021.

Cecilia R. Caglio, R. Matthew Darst, and Şebnem Kalemli-Özcan. Collateral heterogeneity and monetary policy transmission: Evidence from loans to SMEs and large firms. *NBER Working Paper 28685*, 2021.

Dario Caldara and Edward Herbst. Monetary policy, real activity, and credit spreads: Evidence from Bayesian proxy SVARs. *American Economic Journal: Macroeconomics*, 11(1):157–192, 2019.

Anna Cieslak. Short-rate expectations and unexpected returns in Treasury bonds. *Review of Financial Studies*, 31(9):3265–3306, 2018.

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

John H. Cochrane and Monika Piazzesi. The Fed and interest rates—A high-frequency identification. *American Economic Review*, 92(2):90–95, 2002.

Nicolas Crouzet. Credit disintermediation and monetary policy. *IMF Economic Review*, 69(1): 23–89, 2021.

Dominik Damast, Christian Kubitza, and Jakob Ahm Sørensen. Homeowners insurance and the transmission of monetary policy. *SSRN 5119139*, 2025.

Jens Dick-Nielsen. Liquidity biases in TRACE. *Journal of Fixed Income*, 19(2):43–55, 2009.

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

Robin Döttling and Lev Ratnovski. Monetary policy and intangible investment. *Journal of Monetary Economics*, 134:53–72, 2023.

Thomas Drechsel, Daniel Lewis, Davide Melcangi, and Laura Pilossoph. The investment channel of monetary policy: Disentangling firm heterogeneity. *Mimeo*, 2026.

Andrea L. Eisfeldt and Dimitris Papanikolaou. Organization capital and the cross-section of expected returns. *Journal of Finance*, 68(4):1365–1406, 2013.

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

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

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

Jon Faust, John H. Rogers, Eric T. Swanson, and Jonathan H. Wright. Identifying the effects of monetary policy shocks on exchange rates using high frequency data. *Journal of the European Economic Association*, 1(5):1031–1057, 2003.

Jon Faust, Eric T. Swanson, and Jonathan H. Wright. Identifying VARs based on high frequency futures data. *Journal of Monetary Economics*, 51(6):1107–1131, 2004.

Thiago R. T. Ferreira, Daniel A. Ostry, and John Rogers. Firm financial conditions and the transmission of monetary policy. *Finance and Economics Discussion Series 2023-037*, 2023.

William L. Gamber. Firm dynamics, inflation, and the transmission of monetary policy. *Finance and Economics Discussion Series 2026-003*, 2026.

Mark Gertler and Peter Karadi. Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76, 2015.

Matthias Gnewuch and Donghai Zhang. Monetary policy, firm heterogeneity, and the distribution of investment rates. *Journal of Monetary Economics*, 149:103721, 2025.

Roberto Gómez-Cram and Marco Grotteria. Real-time price discovery via verbal communication: Method and application to Fedspeak. *Journal of Financial Economics*, 143(3):993–1025, 2022.

Yuriy Gorodnichenko and Michael Weber. Are sticky prices costly? Evidence from the stock market. *American Economic Review*, 106(1):165–199, 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. *International Journal of Central Banking*, 1(1):55–93, 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.

Alex Hsu, Indrajit Mitra, Yu Xu, and Linghang Zeng. The Fed information effect and firm investment. *Federal Reserve Bank of Atlanta Working Paper 2023-06a*, 2023.

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.

Marek Jarociński and Peter Karadi. Deconstructing monetary policy surprises—The role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43, 2020.

Priit Jeenas. Firm balance sheet liquidity, monetary policy shocks, and investment dynamics. *Mimeo*, 2026.

Priit Jeenas and Ricardo Lagos. Q–monetary transmission. *Journal of Political Economy*, 132(3): 971–1012, 2024.

Joachim Jungherr, Matthias Meier, Timo Reinelt, and Immo Schott. Corporate debt maturity matters for monetary policy. *Federal Reserve Bank of San Francisco Working Paper 2024-30*, 2024.

Şebnem Kalemli-Özcan. U.S. monetary policy and international risk spillovers. *NBER Working Paper 26297*, 2019.

Thomas Kroen, Ernest Liu, Atif R. Mian, and Amir Sufi. Falling rates and rising superstars. *Mimeo*, 2025.

Kenneth N. Kuttner. Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544, 2001.

Aeimit Lakdawala. Decomposing the effects of monetary policy using an external instruments SVAR. *Journal of Applied Econometrics*, 34(6):934–950, 2019.

Aeimit K. Lakdawala and Timothy Moreland. Firm-level uncertainty and the transmission of monetary policy. *Review of Economics and Statistics*, 108(3):807–816, 2026.

Erik Loualiche, Alexandre R. Pecora, Fabricius Somogyi, and Colin Ward. Monetary policy transmission through the exchange rate factor structure. *Journal of Financial Economics*, 182:104305, 2026.

Kurt G. Lunsford. Policy language and information effects in the early days of Federal Reserve forward guidance. *American Economic Review*, 110(9):2899–2934, 2020.

Matthew A. Masten and Alexandre Poirier. The effect of omitted variables on the sign of regression coefficients. *American Economic Review*, 116(7):2685–2710, 2026.

Silvia Miranda-Agrippino. Unsurprising shocks: Information, premia, and the monetary transmission. *Bank of England Staff Working Paper 626*, 2016.

Silvia Miranda-Agrippino and Hélène Rey. U.S. monetary policy and the global financial cycle. *Review of Economic Studies*, 87(6):2754–2776, 2020.

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.

Namrata Narain and Kunal Sangani. The market impact of Fed communications: The role of the press conference. *International Journal of Central Banking*, 22(1):313–389, 2026.

Ricardo Nunes, Ali K. Ozdagli, and Jenny Tang. Interest rate surprises: A tale of two shocks. *Federal Reserve Bank of Boston Working Paper 22-2*, 2022.

Emily Oster. Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204, 2019.

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.

Pascal Paul. The time-varying effect of monetary policy on asset prices. *Review of Economics and Statistics*, 102(4):690–704, 2020.

Altan Pazarbaşı. Cash heterogeneity and the payout channel of monetary policy. *Journal of Finance*, 2026.

Ander Perez-Orive, Yannick Timmer, and Alejandro van der Ghote. Monetary policy under multiple financing constraints. *Finance and Economics Discussion Series 2026-021*, 2026.

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

Valerie A. Ramey. Macroeconomic shocks and their propagation. *Handbook of Macroeconomics*, 2:71–162, 2016.

Christina D. Romer and David H. Romer. A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084, 2004.

Julia Selgrad and Kerry Y. Siani. Monetary policy and investment plans. *SSRN 5342848*, 2026.

James H. Stock and Mark W. Watson. Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *Economic Journal*, 128(610):917–948, 2018.

Eric T. Swanson and Vishuddhi Jayawickrema. Speeches by the Fed chair are more important than FOMC announcements: An improved high-frequency measure of U.S. monetary policy shocks. *Mimeo*, 2024.

Johannes F. Wieland and Mu-Jeung Yang. Financial dampening. *Journal of Money, Credit and Banking*, 52(1):79–113, 2020.

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

A 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.

B Additional Tables

Table A.1: Description of Monetary Policy Surprise Series

Shock	Sample	Notes
Raw shocks:		
Acosta et al. [2025] statement ($stmt$)	1994Q1–2023Q4	
Acosta et al. [2025] monetary event (me)	1994Q1–2023Q4	
Bauer and Swanson [2023] (mps)	1988Q1–2023Q4	
Bu et al. [2021]	1994Q1–2020Q1	
Gorodnichenko and Weber [2016]	1990Q1–2009Q4	
Gürkaynak et al. [2005] target	1995Q1–2023Q4	Extended by Acosta et al. [2024]
Gürkaynak et al. [2005] path	1995Q1–2023Q4	Extended by Acosta et al. [2024]
Nakamura and Steinsson [2018]	1995Q1–2023Q4	Extended by Acosta et al. [2024]
Romer and Romer [2004]	1988Q1–2007Q4	Extended by Wieland and Yang [2020]
Orthogonalized shocks:		
Acosta et al. [2025] statement ($stmt^{\perp}$)	1994Q1–2023Q4	
Acosta et al. [2025] monetary event ($me^{\perp}$)	1994Q1–2023Q4	
Bauer and Swanson [2023] ($mps^{\perp}$)	1988Q1–2023Q4	
Miranda-Agrippino and Ricco [2021] (IV1)	1990Q1–2009Q4	

Note: This table provides further details on the monetary policy surprise series used in Tables 2 and 3. The sample start date for each series is based on data availability (except for Romer and Romer [2004], which is cut off in 1988Q1 to align with Bauer and Swanson [2023]), and the sample end date is set to 2023Q4 unless the series ends earlier. Across all series, we again exclude the post-9/11 meeting and the COVID-19 period.

Table A.2: Return-Sensitivity (β) Estimation: Summary Statistics

	N	Mean	$\mathbb{E}_w[\cdot]$	S.D.	Min	P25	Median	P75	Max
Panel A. Estimated return sensitivity									
$\widehat{\beta}_i$	2,517	-0.21	-0.16	0.17	-1.08	-0.32	-0.21	-0.11	0.54
Panel B. Regression variables									
Two-day return R_{it} (%)	322,799	0.16		4.35	-27.38	-1.87	0.02	2.02	31.75
Policy surprise $stmt_t$	170	0.003		0.030	-0.148	-0.008	0.005	0.016	0.090
Panel C. Sample and W-portfolio composition									
Sample	2,517 firms, 170 FOMC events, 2002–2023 (avg. 128 events/firm)								
Firms with $\widehat{\beta}_i < 0$	91% of firms, 97% of W-portfolio value								
End-2023 W-portfolio value	\$32.8 trillion								
Equity	88% of W-portfolio value								
Equity + traded bonds	95% of W-portfolio value								
+ within-firm imputed bonds	99% of W-portfolio value								
Coverage of Compustat universe	86.6% of non-financial, non-utility universe W value (2,517 of 6,663 firms)								

Note: $\widehat{\beta}_i$ is the slope from regressing firm i's two-day W-portfolio return on the monetary policy surprise (Equation (4)), winsorized at the median $\pm 5\times$IQR, on the baseline sample. Mean, S.D., and percentiles are across firms; $\mathbb{E}_w[\cdot]$ is the value-weighted mean (the aggregation input). Value shares are each firm's average share of the sample's equity-plus-bond value across events.

Table A.3: Investment-Q (γ) Estimation: Summary Statistics

	N	Mean	$\mathbb{E}_w[\cdot]$	S.D.	Min	P25	Median	P75	Max
Panel A. Estimated investment-Q sensitivity									
$\widehat{\gamma}_i$, standard investment	2,517	0.028	0.017	0.041	-0.131	0.006	0.017	0.036	0.165
$\widehat{\gamma}_i$, total investment	2,517	0.048	0.031	0.070	-0.259	0.013	0.039	0.072	0.336
Panel B. Regression variables									
Investment rate I/K, standard	51,697	0.13		0.13	-0.04	0.05	0.09	0.16	0.76
Investment rate I/K, total	51,697	0.17		0.13	-0.02	0.09	0.13	0.20	0.88
Tobin's Q, standard	51,697	3.7		4.8	-14.3	0.7	1.7	4.8	16.7
Tobin's Q, total	51,697	1.2		1.4	-4.5	0.4	0.8	1.5	5.9
Panel C. Sample									
Sample	2,517 firms, 51,697 firm-years, 1994–2023, excl. 2020 (avg. 21 years/firm)								
Firms with $\widehat{\gamma}_i > 0$, standard	88% of firms, 82% of W-portfolio value								
Firms with $\widehat{\gamma}_i > 0$, total	85% of firms, 78% of W-portfolio value								

Note: $\widehat{\gamma}_i$ is the firm-specific slope from the investment-Q sensitivity regression (Equation (5)) with firm and industry$\times$year fixed effects, winsorized at the median $\pm 5\times$IQR. Standard investment uses tangible capital; total investment adds intangibles following Peters and Taylor [2017]. $\mathbb{E}_w[\cdot]$ is the value-weighted mean; I/K and Q are the regression variables.

Table A.4: Quantification of the Aggregate Investment Response: Robustness

#	Specification	Investment-Q sensitivity (γ) $\mathbb{E}_w[\widehat{\gamma}_i]$	N_γ	Return sensitivity (β) $\mathbb{E}_w[\widehat{\beta}_i]$	N_β	Covariance $\text{Cov}_w(\widehat{\gamma}_i, \widehat{\beta}_i)$	Total response % W	% GDP	% K
Baseline specification									
1	Standard investment (baseline)	0.017	2,517	-0.16	2,517	-0.00085	-0.36%	-0.33%	-0.93%
2	Total investment (+intangibles)	0.031	2,517	-0.16	2,517	-0.00075	-0.58%	-0.53%	-0.76%
Panel A. Controls in the return-sensitivity regression									
3	8 lags of 1-year Treasury yield changes	0.017	2,517	-0.17	2,517	-0.00078	-0.36%	-0.33%	-0.94%
4	4 lags of GDP, inflation, unemp.	0.017	2,517	-0.16	2,517	-0.00076	-0.35%	-0.32%	-0.90%
5	Bauer-Swanson controls	0.017	2,517	-0.18	2,517	-0.00062	-0.36%	-0.33%	-0.93%
6	All controls	0.017	2,517	-0.18	2,517	-0.00076	-0.38%	-0.35%	-1.00%
Panel B. Alternative return-sensitivity measurement									
7	Monetary-event surprise (me)	0.017	2,517	-0.15	2,517	-0.00066	-0.32%	-0.29%	-0.83%
8	Statement surprise, excl. GFC	0.017	2,517	-0.17	2,517	-0.00077	-0.37%	-0.34%	-0.95%
9	One-day return window	0.017	2,517	-0.08	2,517	-0.00027	-0.15%	-0.14%	-0.40%
10	Equity only, longer sample (1994–2023)	0.016	3,018	-0.17	3,018	-0.00082	-0.37%	-0.29%	-1.65%
Panel C. Treatment of firms with untraded bonds									
11	Impute firm's own bond return only	0.017	2,517	-0.17	2,517	-0.00092	-0.37%	-0.34%	-0.94%
12	Actually-traded bonds only	0.016	2,517	-0.17	2,517	-0.00099	-0.37%	-0.33%	-0.93%
Panel D. Fixed effects in the investment-Q sensitivity regression									
13	Only firm and time FEs	0.021	2,517	-0.16	2,517	-0.00090	-0.43%	-0.40%	-1.12%
14	Drop time & industry $\times$ time FEs	0.023	2,517	-0.16	2,517	-0.00085	-0.46%	-0.43%	-1.20%
Panel E. Alternative unit of analysis for estimating the sensitivities									
15	Size $\times$ industry cells	0.012	2,517	-0.18	2,517	-0.00026	-0.24%	-0.22%	-0.63%
Panel F. Relaxing the sample restrictions									
16	Fewer FOMC meetings (25+)	0.017	2,857	-0.16	2,857	-0.00078	-0.35%	-0.33%	-0.91%
17	Fewer investment-Q years (5+)	0.017	3,107	-0.16	3,107	-0.00090	-0.37%	-0.35%	-0.94%
18	β & γ on own sample restrictions	0.017	3,017	-0.17	3,282	-0.00088	-0.37%	-0.38%	-0.99%
Panel G. Other: winsorization, shrinkage, sub-periods, estimation error									
19	Coefficient winsorization at 3$\times$ IQR	0.016	2,517	-0.16	2,517	-0.00077	-0.33%	-0.31%	-0.87%
20	Shrink both (empirical Bayes)	0.014	2,517	-0.16	2,517	-0.00046	-0.28%	-0.25%	-0.71%
21	Winsorize firm value weights	0.023	2,517	-0.20	2,517	-0.00085	-0.54%	-0.50%	-1.39%
22	Shrink both + winsorize value weights	0.018	2,517	-0.19	2,517	-0.00043	-0.38%	-0.36%	-1.00%
23	Early half of sample	0.026	1,556	-0.15	1,556	-0.00026	-0.41%	-0.29%	-1.01%
24	Late half of sample	0.015	1,482	-0.17	1,482	-0.00054	-0.31%	-0.26%	-0.80%
25	Covariance adjusted for estimation error	0.017	2,517	-0.16	2,517	-0.00080	-0.35%	-0.33%	-0.92%

Note: Each row re-estimates the implied aggregate investment response $\mathbb{E}_w[\gamma] \cdot \mathbb{E}_w[\beta] + \text{Cov}_w(\gamma, \beta)$, varying only the ingredient named in the row or panel heading and holding the rest at baseline (row 1, standard investment). The "Total response" columns express it as a share of the covered portfolio W (% W) and rescaled by W/GDP and W/K on each row's own sample. N_γ and N_β are the γ- and β-sample firm counts, equal except in the row where β and γ are each estimated on their own sample restrictions. See Section 4.6 for further description of each row.

C Additional Figures

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

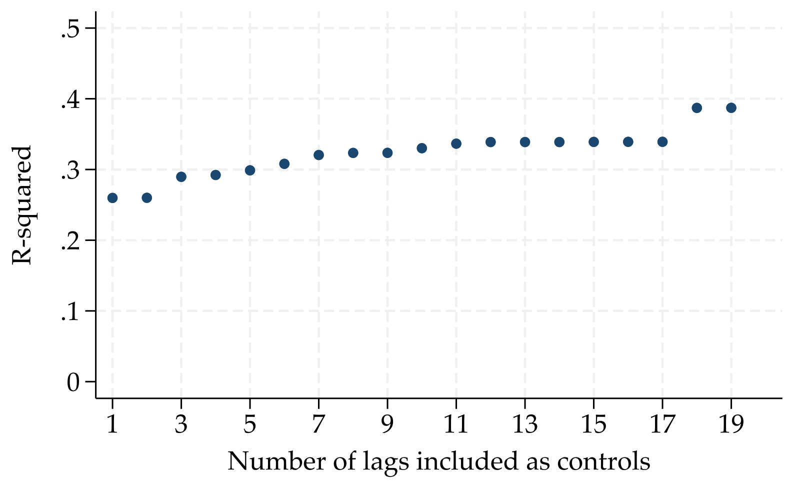

Note: We estimate the predictive regressions: $stmt_t = \alpha + \sum_{l=1}^{k} \beta_l (y_{t-l}^{(1)} - y_{t-l-1}^{(1)}) + \epsilon_t$ using various lags up to $k = 19$. $stmt_t$ is the monetary policy surprise in quarter t, which is the sum of all the monetary policy statement surprises from Acosta 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: Path of One-Year Yield Around Monetary Policy Statement Surprise: Monthly Frequency

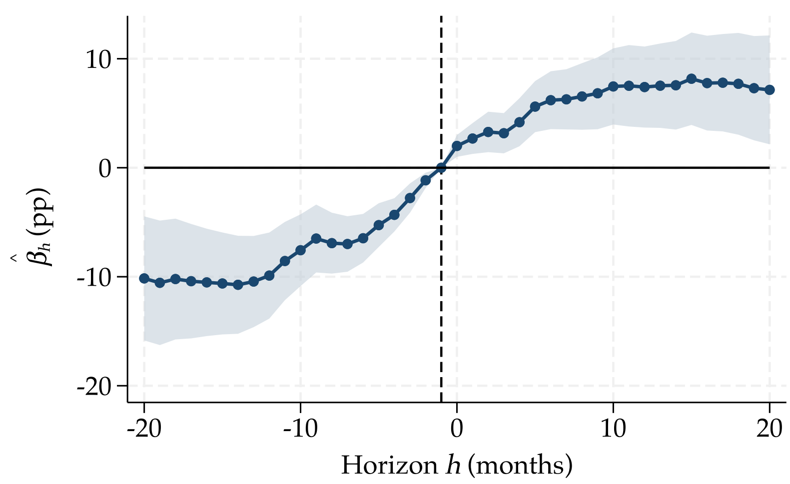

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-20}^{20}$ estimated from Equation (1). Monetary policy statement surprises come from Acosta et al. [2025]. The outcome variable is the change in the one-year $(y_t^{(1)})$ Treasury zero-coupon bond yield. Confidence intervals are constructed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

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

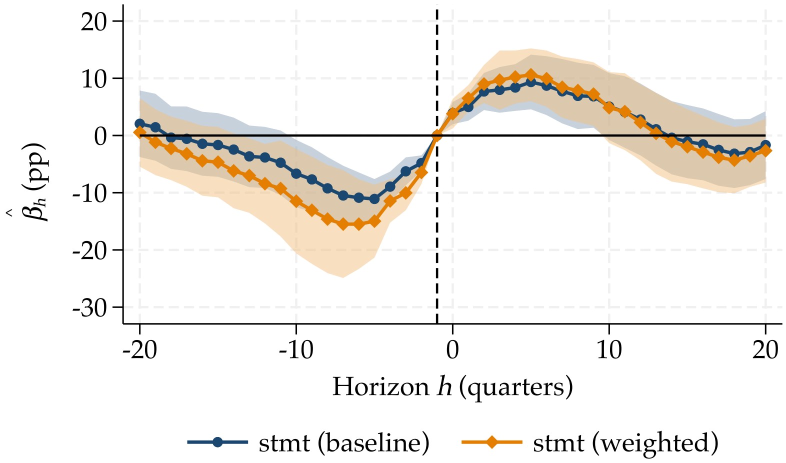

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-20}^{20}$ estimated from Equation (1). Monetary policy statement surprises come from Acosta et al. [2025]. They are aggregated within a quarter using either an unweighted sum, or a weighted-average using the fraction of the quarter remaining 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.4: Monetary Policy *Event* Surprises and Interest Rates

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

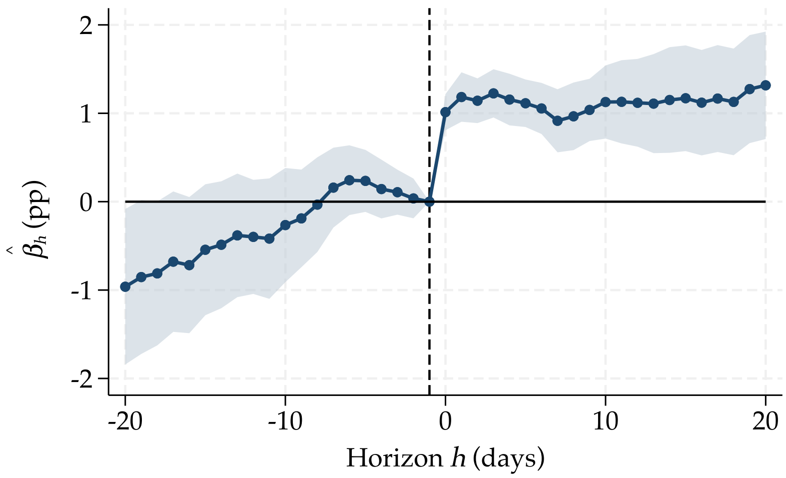

(b) $y^{(1)}_{t+h} - y^{(1)}_{t-1} = \hat{\beta}_h \cdot me_t$, Quarterly Frequency

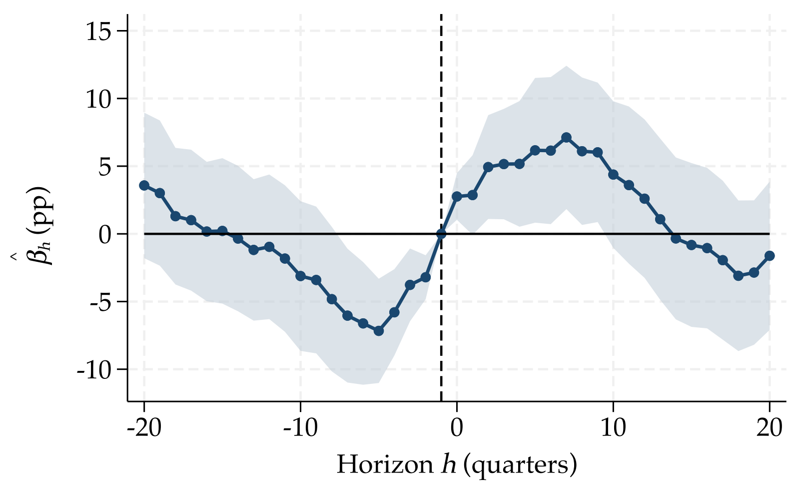

Note: This figure plots the sequence $\{\hat{\beta}_h\}_{h=-20}^{20}$ estimated from Equation (1) where we use the monetary policy *event* surprises from Acosta et al. [2025]. Panel (a) uses daily data and daily surprises; Panel (b) uses quarterly data and quarterly-aggregated surprises. The outcome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

Figure A.5: Return of Aggregate Portfolio of Bonds and Stocks Around Monetary Policy Surprises

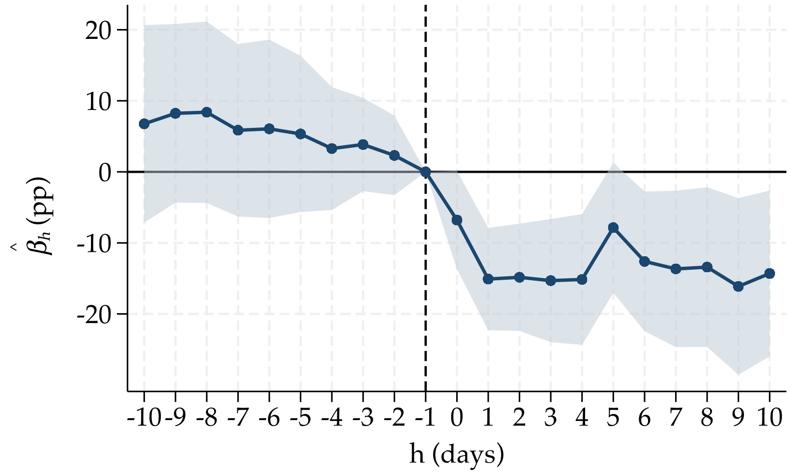

Note: Sample period 2002–2023. We plot the sequence $\{\hat{\beta}_h\}_{h=-10}^{10}$ estimated from the per-horizon local projection:

$$R_{t-1\to t+h} = \alpha_h + \beta_h \cdot stmt_t + \epsilon_{t,h},$$

where $R_{t-1\to t+h}$ is the cumulative return on the value-weighted portfolio of bonds and stocks in our sample—the aggregate W portfolio introduced in Section 3—measured between the close of day $t-1$ and the close of day $t+h$, normalized to zero on day $t-1$. Monetary policy statement surprises, $stmt_t$, come from Acosta et al. [2025]. Confidence intervals are constructed at the 90% level with heteroskedasticity-robust standard errors.

Figure A.6: Stock Returns Around Monetary Policy Statement Surprises

(a) Baseline Sample (2002-2023)

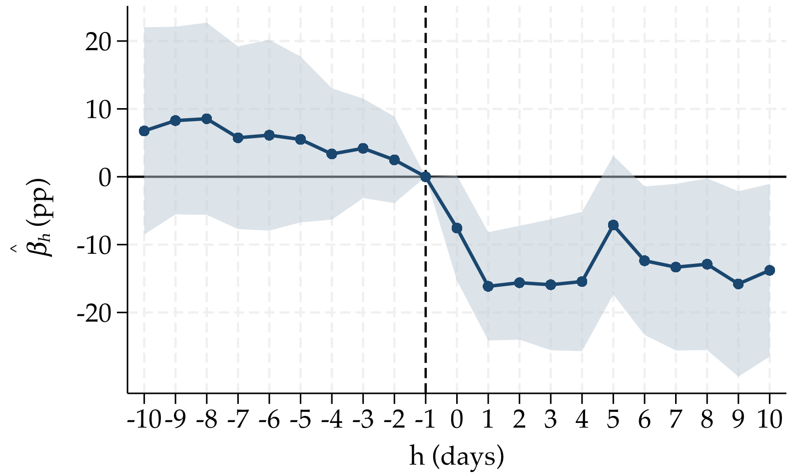

(b) Longer Sample (1994-2023)

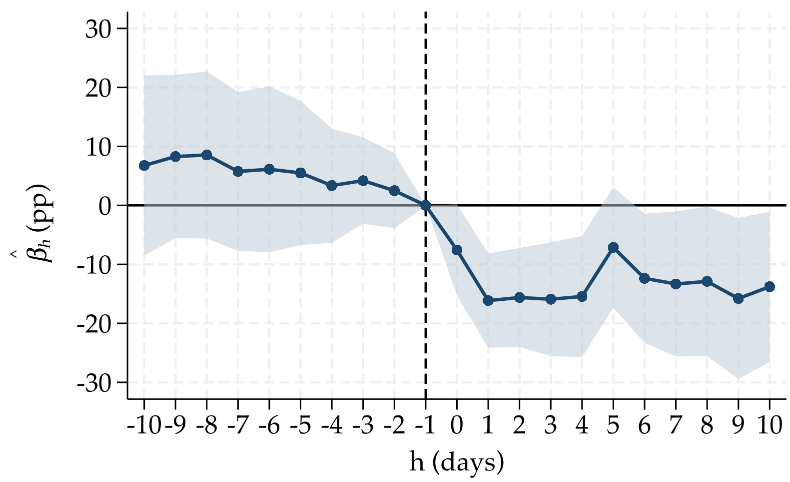

Note: This figure shows cumulative aggregate stock returns around FOMC meetings in response to monetary policy statement surprises. Returns are computed for a value-weighted portfolio of equities issued by firms in our sample. Panel (a) uses the baseline sample period 2002–2023, reflecting the availability 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 normalized to zero on day $t-1$ and plotted from $t-10$ to $t+10$, estimated by the per-horizon local projection of Figure A.5. 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. Confidence intervals are constructed at the 90% level with heteroskedasticity-robust standard errors.

Figure A.7: Bond Returns Around Monetary Policy Statement Surprises

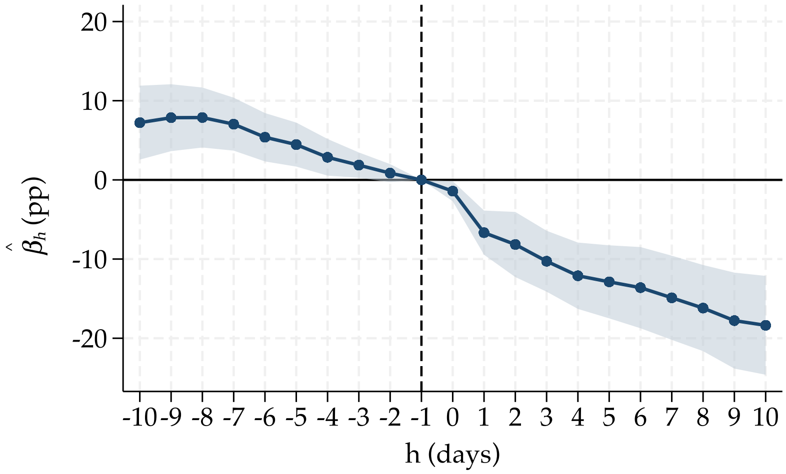

Note: This figure shows cumulative returns on the aggregate value-weighted portfolio of corporate bonds around FOMC meetings in response to monetary policy statement surprises. Bond returns are computed using TRACE transaction data by constructing bond-day prices as volume-weighted averages of clean trades. For each meeting at date t, cumulative returns are normalized to zero on day $t-1$ and plotted from $t-10$ to $t+10$, estimated by the per-horizon local projection of Figure A.5. 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. Confidence intervals are constructed at the 90% level with heteroskedasticity-robust standard errors.

Figure A.8: Return of Aggregate Portfolio of Bonds and Stocks Around Monetary Policy Surprises: Adding Controls

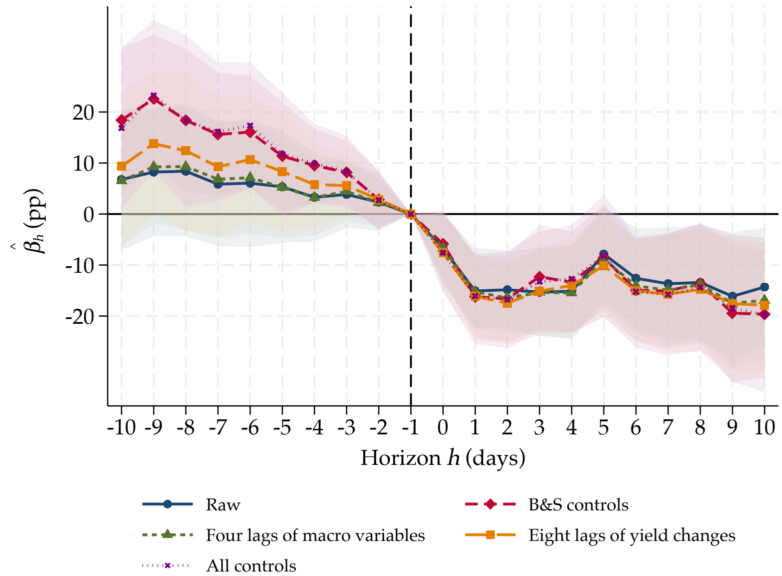

Note: This figure overlays the local projection of Figure A.5 (the cumulative return on the aggregate W portfolio of bonds and stocks, normalized to zero on day $t-1$) re-estimated with each set of low-frequency controls used in Figure 5: no controls (raw), the Bauer and Swanson [2023] controls, four lags of the macro variables, eight lags of yield changes, and all controls together; controls enter linearly at each horizon. Sample period 2002–2023. Monetary policy statement surprises, $stmt_t$, come from Acosta et al. [2025]. Confidence intervals are constructed at the 90% level with heteroskedasticity-robust standard errors.

Figure A.9: Cross-Validation of Return Sensitivity Estimates

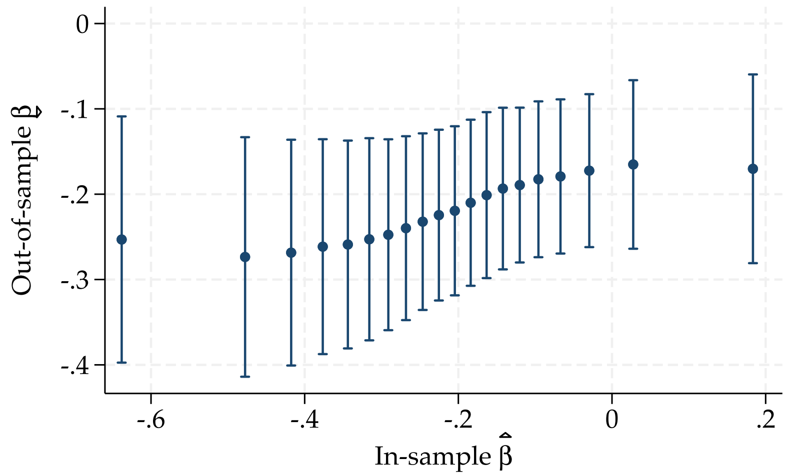

Note: This figure reports a repeated cross-validation of our firm-level return sensitivities. In each repetition, we split the FOMC events into four folds—assigned so that each fold spans the full range of surprises—and cycle through them so that every event is held out for testing exactly once, each time training on three folds and testing on the remaining one (a 75/25 split). On the training folds we estimate raw firm-level slopes $\widehat{\beta}_i$ from Equation (4), winsorize them at the median $\pm\, 5\times$IQR, sort firms into 20 equal-sized bins by these slopes, and record each bin's average in-sample $\widehat{\beta}_i$; we then set aside the held-out observations, each tagged with its firm's bin. Stacking the held-out observations from all four folds, we estimate one out-of-sample slope per bin from a pooled regression of returns on the surprise interacted with bin indicators, absorbing fold$\times$firm fixed effects. We repeat this over 50 random fold assignments and, within each bin, average the in-sample $\widehat{\beta}_i$ and the out-of-sample slope. The x-axis reports the average in-sample $\widehat{\beta}_i$ in each bin and the y-axis the corresponding out-of-sample slope; error bars are 95% confidence intervals from standard errors clustered by firm and event.

Figure A.10: Cross-Validation of Investment-Q Sensitivity Estimates

(a) Standard Investment (b) Total Investment

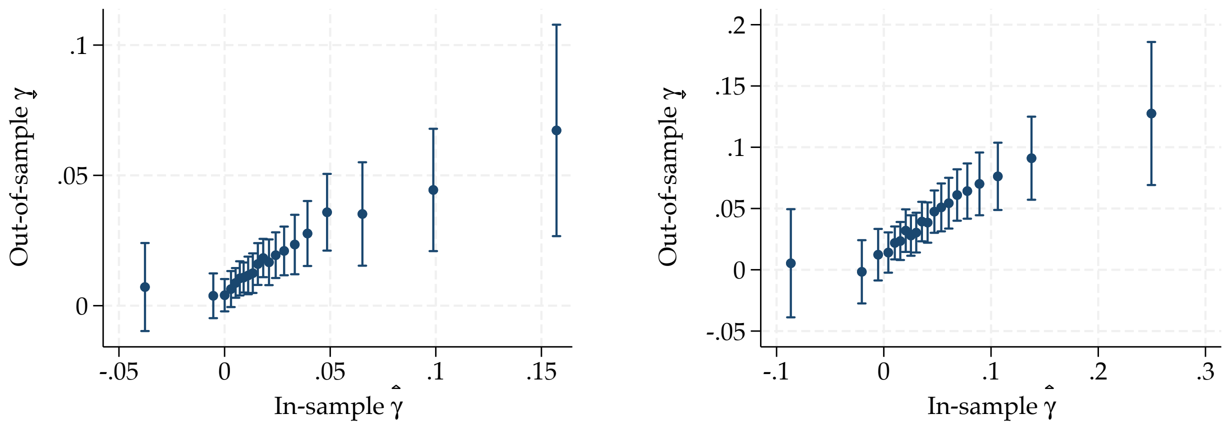

Note: This figure reports a leave-one-year-out cross-validation of our firm-level investment-Q sensitivities. For each year in turn, we hold out that year's observations, estimate raw firm-level slopes $\widehat{\gamma}_i$ from Equation (5) on the remaining years, winsorize them at the median $\pm\, 5\times$IQR, and sort firms into 20 equal-sized bins by these slopes. In the held-out year we then estimate one out-of-sample slope per bin from a within-bin regression of investment on lagged Q, both purged of firm and industry×year fixed effects. Cycling through all years so that each is held out once, we average the in-sample $\widehat{\gamma}_i$ and the out-of-sample slope within each bin across the held-out years. Panel (a) uses the standard (tangible) investment measure and panel (b) the total-investment measure of Peters and Taylor [2017]. The x-axis reports the average in-sample $\widehat{\gamma}_i$ in each bin and the y-axis the corresponding out-of-sample slope; error bars are 95% confidence intervals from heteroskedasticity-robust standard errors, averaged across the held-out years.

Figure A.11: Out-of-Sample Investment Predictability: Homogeneous-γ Comparison

(a) Standard Investment

(b) Total Investment

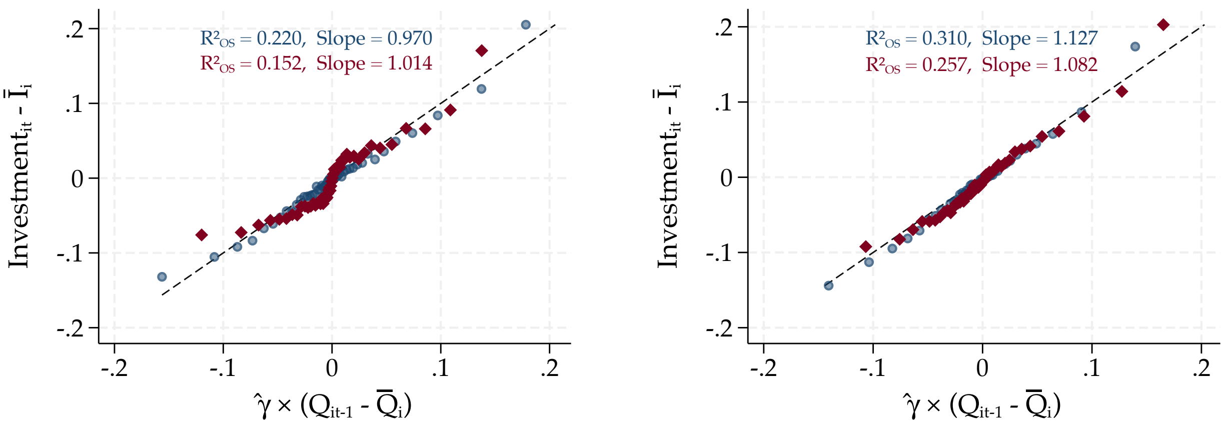

Note: Each panel plots binned averages of actual demeaned investment (y-axis) against predicted demeaned investment (x-axis). Prediction uses training-sample $\widehat{\gamma}_i$ estimates, and firms' investment and Q are demeaned using training-sample means. The dashed 45-degree line marks perfect prediction (actual equals predicted). Each panel overlays two sets of binned predictions: our baseline *heterogeneous*-γ model, which assigns every firm its own $\widehat{\gamma}_i$, and a restricted *homogeneous*-γ benchmark that imposes a single common slope on all firms. The in-panel labels report each model's out-of-sample R^2 and the slope of actual on predicted investment.

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.[57] On the assets side, in our Compustat sample, about 35% 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 31% of current assets and roughly 56 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 24% of current assets, and also have short effective duration because firms replenish them frequently: in our sample, Days Inventory Outstanding is about 74 days. 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.[58] 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.

[57] These working-capital statistics are computed on the firms in our analysis sample over fiscal years 2002–2023, restricted to firm-years reporting a complete current-account breakdown. The composition shares are dollar-weighted (aggregate ratios computed each year, then averaged across firm-years), whereas the turnover measures (Days Sales, Inventory, and Payable Outstanding) are firm-level ratios winsorized at the median $\pm 5\times$IQR.

[58] In our sample, the average Days Payable Outstanding is about 51.