Climate Policy Uncertainty and Financial Transmission: Evidence from U.S. Credit Markets*

Facundo Luna
Rutgers University

This version: March 2026

Abstract

Climate policy uncertainty (CPU) has surged to record levels, yet its transmission through financial markets remains largely unexplored. Using the narrative identification strategy of Gavriilidis et al. (2026), I document three key findings on the credit market transmission of CPU shocks. First, CPU shocks significantly tighten credit markets: excess bond premiums rise, and bank loans contract by 1–2%, driven initially by stricter bank lending standards. Second, these effects are strongly state-dependent, with investment declines and credit spread widenings severely amplified when aggregate financial conditions are already tight. Third, firm-level data reveals that climate-exposed and highly leveraged firms face the steepest rises in borrowing costs; however, less-leveraged firms cut investment more sharply, reflecting greater financial flexibility to adjust (Ottonello and Winberry, 2020). These findings highlight a novel financial amplification channel with direct implications for climate stress-testing and transition policy design.

JEL Classification: E44, G21, H23, Q54, Q58
Keywords: Climate policy uncertainty, financial transmission, credit markets, local projections, state dependence, excess bond premium

*Corresponding author: Facundo Luna, Department of Economics, Rutgers University. Email: facundo.luna@rutgers.edu. Website: facundolunam.github.io.

1 Introduction

Climate policy in the United States has been characterized by increasingly large swings: from cap-and-trade legislation under Obama to deregulation under Trump, to the Inflation Reduction Act under Biden, and most recently to renewed fossil-fuel promotion under the second Trump administration. These reversals have pushed climate policy uncertainty to record levels (Gavriilidis et al., 2026; Gavriilidis, 2021). A growing body of work documents that this uncertainty carries real macroeconomic costs—reducing output, investment, and employment while simultaneously raising prices (Gavriilidis et al., 2026; Luna-Mallea, 2026). Yet despite the centrality of financial frictions in the broader uncertainty literature (Bloom, 2009; Gilchrist and Zakrajšek, 2012; Caldara et al., 2016), the financial transmission of climate policy uncertainty remains entirely unexplored.

This gap is striking. The uncertainty literature identifies credit markets as a key amplification mechanism: uncertainty shocks raise risk premia, tighten bank lending standards, and contract credit supply, magnifying the real effects of the initial shock (Gilchrist and Zakrajšek, 2012; Caldara et al., 2016; Baker et al., 2016). Luna-Mallea (2026) proposes, in a theoretical DSGE framework, that CPU transmits through an asset-quality channel—brown-capital collateral values decline under policy uncertainty, tightening credit constraints and amplifying downturns—but provides only aggregate macroeconomic evidence without directly testing the financial mechanism. Gavriilidis et al. (2026) provide definitive causal identification of CPU's macroeconomic effects using a novel narrative instrument, but their analysis contains no financial variables beyond the short-term interest rate. Neither paper—nor any other in the literature—examines whether CPU shocks transmit through credit spreads, bank lending, or corporate bond markets.

This paper bridges this gap. Using the Gavriilidis et al. (2026) narrative identification strategy—146 exogenous climate policy events over 1986–2019, yielding a first-stage F-statistic of approximately 23.5—I provide new evidence on the transmission of CPU shocks through U.S. credit markets. The analysis proceeds in three steps.

First, I estimate linear local projections (Jordà, 2005) for a broad set of financial variables with lack of evidence in the CPU context: the Gilchrist and Zakrajšek (2012) excess bond premium (EBP) and GZ credit spread, commercial and industrial (C&I) bank loans, the Senior Loan Officer Opinion Survey (SLOOS) on lending standards, ICE BofA corporate credit spreads by rating (AAA, BBB, high yield), and the S&P 500.

Second, I examine whether these financial effects are state-dependent on prevailing financial conditions. Following the smooth-transition approach of Auerbach and Gorodnichenko (2012) and Tenreyro and Thwaites (2016), I interact the CPU shock with measures of financial tightness—the Chicago Fed's Adjusted National Financial

Conditions Index (ANFCI), the excess bond premium, and the GZ credit spread—to test whether CPU's financial impact is amplified when the financial system is already stressed.

Third, I move to the firm level, exploiting the Compustat panel to test whether financially constrained firms experience disproportionately larger investment and R&D declines following CPU shocks, following the heterogeneity approach in Gavriilidis et al. (2026) and Ottonello and Winberry (2020).

The results reveal a powerful and previously undocumented financial transmission channel. First, at the aggregate level, a one-standard-deviation CPU shock raises the GZ credit spread by approximately 5–8 basis points, increases the excess bond premium by a similar magnitude, and widens high-yield corporate spreads by roughly 15–25 basis points over 12–24 months. The shock also produces a persistent contraction in commercial and industrial bank lending on the order of 1–2%, with the Senior Loan Officer Opinion Survey confirming that banks actively tighten lending standards in the first one to four quarters following the shock. A decomposition exercise shows that conditioning on observed lending standards absorbs the short-run credit contraction, while the longer-run decline reflects additional general equilibrium feedback and unmeasured dimensions of credit rationing.

Second, these financial effects are strongly state-dependent. When the Adjusted National Financial Conditions Index indicates tight conditions, CPU shocks generate investment declines several times larger than during loose conditions, accompanied by sharp increases in the excess bond premium and high-yield spreads that are essentially absent in the loose regime. The credit supply channel is similarly asymmetric: banks tighten lending standards substantially for small firms in tight conditions, with equality of responses across states rejected at conventional levels.

Third, firm-level evidence from Compustat and WRDS corporate bond data reveals meaningful heterogeneity. Firms with greater climate change exposure (Sautner et al., 2023) experience larger investment declines following CPU shocks. In the corporate bond market, higher-leverage firms face rising bond yields and declining bond volumes, indicating that the financial channel penalizes leveraged balance sheets. However, the Compustat investment-leverage interaction reveals a more nuanced pattern—consistent with Ottonello and Winberry (2020)—in which less-leveraged firms exhibit larger investment cuts, suggesting that financial flexibility enables sharper adjustment to uncertainty rather than that constraints amplify the response.

This paper makes three contributions to the literature. First, I provide the first direct empirical evidence on the financial transmission of narratively identified CPU shocks—filling a major gap in Gavriilidis et al. (2026), whose comprehensive macro analysis includes no financial market variables. The finding that CPU shocks contract bank lending by 1–2% is, to my knowledge, new to the literature.

Second, I document that the financial effects of CPU are strongly state-dependent on financial conditions. When the ANFCI indicates tight financial conditions, CPU shocks raise the excess bond premium significantly, widen high-yield and BBB corporate spreads, and generate equity declines several times larger than during loose conditions. This financial-conditions amplification differs fundamentally from the business-cycle state dependence documented in Luna-Mallea (2026), which finds larger effects during expansions. The financial-conditions channel reveals a "doom loop": CPU shocks are most damaging precisely when credit markets are already stressed, creating the potential for self-reinforcing cycles of policy uncertainty and financial tightness.

Third, I reconcile two characterizations of CPU in the literature. Gavriilidis et al. (2026) document that CPU transmits as a supply shock (output falls, prices rise), while Luna-Mallea (2026) emphasizes demand-side financial amplification through collateral channels. My results show that both channels coexist: CPU raises risk premia and tightens lending (financial/demand side) while also increasing commodity prices (supply side). The relative importance of each channel depends on the state of financial conditions—when credit is tight, the financial amplification dominates.

These findings are relevant for three policy audiences: central banks, financial regulators, and climate policymakers. For central banks, CPU's financial amplification channel suggests that standard monetary easing may be insufficient to offset the contractionary effects of climate policy uncertainty—particularly when financial conditions are already strained, as the monetary authority faces a simultaneous output-inflation trade-off (Gavriilidis et al., 2026). For financial regulators, the state-dependent credit market responses suggest that NGFS climate stress tests should incorporate climate *policy* uncertainty scenarios, not only physical risk scenarios. For climate policymakers, the results underscore that credible, predictable climate policy is especially valuable when financial conditions are fragile: policy reversals during periods of tight credit coincide with disproportionately large economic damage.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature and positions the contribution. Section 3 describes the data. Section 4 presents the empirical methodology. Section 5 reports the aggregate results. Section 6 presents the firm-level evidence. Section 7 synthesizes the mechanisms. Section 8 draws policy implications. Section 9 concludes.

2 Related Literature

This paper contributes to three strands of literature: the measurement and macroeconomic effects of climate policy uncertainty, the role of financial frictions in propagating uncertainty shocks, and the emerging literature on climate-related financial risk.

2.1 Climate Policy Uncertainty: Measurement and Macroeconomic Effects

The measurement of climate policy uncertainty has advanced rapidly in recent years. Gavriilidis (2021) constructs a news-based Climate Policy Uncertainty (CPU) index for the United States using newspaper coverage of climate-related policy and uncertainty terms. This index has become widely used in the literature, alongside related measures that expand coverage while following a similar news-based methodology (Noailly and Smeets, 2022; Palikhe et al., 2024). In a related approach, Baker et al. (2026) develop the Equity Market Volatility Tracker across several sectors of the economy, including a subindex for Energy and Environmental Regulation. More recently, Gavriilidis et al. (2026) substantially advance the measurement agenda by combining a refined news-based CPU index with a narrative instrument based on 146 exogenous climate policy events, classified by their direction of policy stringency between 1985 and 2019. They identify CPU shocks using a structural VAR with external instruments (SVAR-IV), obtaining a strong first-stage relationship (F-statistic $\approx$ 23.5) and passing standard exogeneity tests relative to financial and macroeconomic shocks.

Empirical evidence shows that climate policy uncertainty has meaningful economic effects. Palikhe et al. (2024) show that higher environmental policy uncertainty reduces investment and employment in environmentally regulated industries, with effects influenced by political cycles. Noailly et al. (2022) find that greater uncertainty lowers the probability that clean technology startups obtain venture capital and increases the stock volatility of green firms. Extending this approach, Noailly et al. (2024) construct a granular Environmental and Climate Policy index that captures the volume, sentiment, and topics of environmental news, showing that increased policy salience supports clean investment while weighing on returns in brown sectors. Beyond the United States, Berestycki et al. (2022) find that climate policy uncertainty reduces investment in pollution-intensive sectors across OECD countries, while Gavriilidis (2021) document that spikes in U.S. climate policy uncertainty coincide with declines in CO_2 emissions, consistent with a reallocation away from carbon-intensive production. At the firm level, Ren et al. (2022) show that climate policy uncertainty reduces productivity in Chinese firms by constraining RD and internal cash flow.

On the aggregate effects side, Gavriilidis et al. (2026) show that CPU shocks reduce industrial production, increase unemployment, raise commodity prices and headline inflation, and suppress firm-level investment and RD. They characterize these dynamics as resembling a negative supply shock, in contrast with the demand-shock properties typically associated with broader economic policy uncertainty shocks (Baker et al., 2016). They also document heterogeneous responses at the firm level, with stronger effects for firms more exposed to climate transition risk (Sautner et al., 2023).

Luna-Mallea (2026) complements this empirical literature with a structural perspective. Using a CPU measure constructed from the Environmental Policy Uncertainty index (Palikhe et al., 2024) and the Energy and Environmental Regulation volatility index from the Equity Market Volatility Tracker (Baker et al., 2026), combined through principal component analysis, the paper estimates local projections and develops a two-sector (green–brown) New Keynesian DSGE model with financial frictions and Epstein–Zin preferences. The results show that CPU shocks are asymmetrically stronger over the business cycle and that the transmission operates primarily through an asset-quality channel: the value of brown capital declines under higher CPU, tightening credit constraints and amplifying downturns. However, the empirical analysis relies mainly on macroeconomic aggregates such as industrial production, unemployment, CPI, and interest rates, and therefore does not directly test the financial transmission mechanism emphasized by the model.

This paper bridges the gap between Gavriilidis et al. (2026) and Luna-Mallea (2026). I use the stronger narrative instrument developed by GKRS to provide direct empirical evidence on the financial transmission channel identified in Luna's DSGE model as the dominant mechanism, but which neither paper directly tests using financial data.

2.2 Uncertainty and Financial Markets

A large macroeconomics literature documents the role of financial markets in propagating uncertainty shocks. Bloom (2009) shows that uncertainty shocks generate rapid drops in output and employment through a "wait-and-see" mechanism that delays hiring and investment. Gilchrist and Zakrajšek (2012) decompose corporate bond credit spreads into an expected-default component and a residual—the excess bond premium (EBP)—and show that shocks to the EBP, which capture fluctuations in the effective risk-bearing capacity of the financial sector, are powerful predictors of economic downturns. Caldara et al. (2016) demonstrate that the real effects of uncertainty are substantially amplified by financial frictions, with credit spreads playing a central role.

The connection between financial conditions and uncertainty transmission is studied in several frameworks. Baker et al. (2016) show that EPU shocks reduce investment and employment, with effects concentrated in policy-sensitive sectors. Ludvigson et al. (2021) distinguish between financial and macroeconomic uncertainty, showing that financial uncertainty shocks are the primary drivers of economic fluctuations. Basu and Bundick (2017) provide a theoretical framework in which uncertainty shocks operate as aggregate demand shocks through a precautionary-savings channel. Ottonello and Winberry (2020) study how monetary policy transmission depends on firm-level financial constraints, showing that low-leverage firms respond more to interest rate changes—a finding that motivates my firm-level heterogeneity analysis.

Despite this rich literature linking uncertainty to financial markets, no paper has examined whether *climate policy* uncertainty—which GKRS show operates as a supply shock rather than a demand shock—transmits through credit markets. My paper provides the first evidence on this question.

2.3 Climate-Related Financial Risk

A growing finance literature studies how climate risks affect financial markets and institutions. Bolton and Kacperczyk (2021) show that carbon emissions are priced in the cross-section of stock returns. Sautner et al. (2023) develop firm-level measures of climate change exposure from earnings calls and show these predict green innovation, green hiring, and are priced in options and equity markets. Ilhan et al. (2021) demonstrate that climate policy uncertainty is priced in option markets, with carbon-intense firms bearing larger downside tail risk premia. Hsu et al. (2023) show that firms in states with higher climate risk earn higher stock returns, consistent with a risk premium.

On the financial intermediation side, Diluiso et al. (2021) study climate transition risk in a macrofinancial model with a banking sector. Carattini et al. (2023) analyze how climate policy affects green bond markets. The Network for Greening the Financial System (NGFS) has developed climate stress-testing scenarios for financial institutions, but these focus primarily on physical and transition risk under known policy paths, not on the effects of policy *uncertainty* on financial conditions.

My paper contributes to this literature by providing causal evidence—based on narratively identified exogenous shocks—that climate policy uncertainty transmits through credit markets. While the existing climate finance literature largely studies cross-sectional pricing of climate risk or asset-level exposures, I focus on the aggregate financial conditions channel: how CPU shocks affect the overall supply and price of credit in the economy, and how these effects depend on prevailing financial conditions.

2.4 State-Dependent Effects of Uncertainty

A smaller literature examines whether the effects of uncertainty vary across economic or financial states. Tenreyro and Thwaites (2016) find that the effects of monetary policy are weaker in recessions, motivating smooth-transition local projection methods. Caggiano et al. (2017) show that uncertainty shocks have larger effects in recessions. Auerbach and Gorodnichenko (2012) introduce the smooth-transition function approach to study fiscal multipliers across the business cycle.

Luna-Mallea (2026) applies state dependence to the CPU context, finding that CPU shocks are more contractionary during expansions than recessions—the opposite of standard uncertainty shocks. My paper advances this by shifting the conditioning variable from the business cycle to financial conditions. Since Luna's own DSGE model attributes CPU's transmission to financial frictions (the asset-quality channel), the natural state variable should be a measure of financial conditions, not the output gap. The ANFCI-based state dependence I document provides a more precise test of the theoretical mechanism.

3 Data

3.1 Climate Policy Uncertainty Shock

The key identifying variation in this paper comes from the climate policy uncertainty shock constructed by Gavriilidis et al. (2026). Their approach proceeds in two steps. First, they construct a monthly news-based CPU index by measuring the frequency of articles in major U.S. newspapers (the *New York Times*, *Wall Street Journal*, *Washington Post*, and *Los Angeles Times*) that jointly discuss climate change, policy, and uncertainty, using carefully curated narrow and broad dictionaries. Second, they identify exogenous variation in CPU using a narrative instrument based on 146 climate policy events between 1985 and 2019, classified by their direction of stringency (tightening vs. loosening). The instrument is used within a structural VAR with external instruments (SVAR-IV) to extract the structural CPU shock.

I use the extracted CPU shock series directly from the Gavriilidis et al. (2026).[1] The shock is standardized (zero mean, unit variance) and spans 1986:M1 to 2019:M12 (408 monthly observations). This is my baseline sample for all aggregate-level analysis.

Table 1 provides summary statistics. The CPU shock has a first-stage F-statistic of approximately 23.5 when the narrative instrument is used, well above the Stock and Yogo (2005) weak-instrument threshold. The narrative instrument is constructed to be orthogonal to contemporaneous macroeconomic and financial shocks (Gavriilidis et al., 2026, Table D.2), which supports the exogeneity assumption.

I use the CPU shock as the baseline regressor in my local projections, following the direct-LP approach that Gavriilidis et al. (2026) emphasize in their main results.

[1] The CPU index and shock data are available at https://github.com/dkaenzig/Climate-Policy-Uncertainty-Index

3.2 Financial Variables

The core contribution of this paper relative to prior work is the examination of financial market variables that have not been studied in the CPU context. I assemble the following data:

Excess bond premium and GZ credit spread. Gilchrist and Zakrajšek (2012) decomposes corporate bond credit spreads into a component attributable to expected default risk and a residual—the excess bond premium (EBP)—that captures fluctuations in the effective risk-bearing capacity of the financial sector. I obtain the monthly EBP and GZ credit spread series from the Federal Reserve Board website.[2] The EBP is a powerful predictor of economic downturns (Favara et al., 2016) and has become a standard measure of credit market sentiment. Both series are available from 1973 through the present.

Bank lending standards. I use the Senior Loan Officer Opinion Survey (SLOOS) on bank lending practices, published quarterly by the Federal Reserve Board.[3] The SLOOS reports the net percentage of banks tightening standards for commercial and industrial loans, providing a direct measure of credit supply conditions. The series for large-firm C&I loans begins in 1990:Q2. For the quarterly SLOOS analysis, I aggregate the monthly CPU shock to the quarterly frequency by summing within each quarter.

Commercial and industrial loans. I measure the volume of bank credit using total C&I loans at all commercial banks (FRED: BUSLOANS), available monthly from 1947. I use the log transformation ($100 \times \ln$) following standard practice, so that impulse responses can be interpreted as cumulative percent changes.

Corporate credit spreads by rating. To examine how CPU affects credit pricing across the risk spectrum, I use the ICE BofA option-adjusted spreads for investment-grade (FRED: BAMLC0A0CM), BBB-rated (FRED: BAMLC0A4CBBB), AAA-rated (FRED: BAMLC0A1CAAA), and high-yield (FRED: BAMLH0A0HYM2) corporate bonds. These daily series, averaged to a monthly frequency, are available from 1996 and measure the spread between corporate bond yields and a matched-maturity Treasury yield curve.

Equity market. I include the S&P 500 index (FRED: SP500) in log levels ($100 \times \ln$) as a benchmark financial variable and to facilitate comparison with Gavriilidis et al.

[2] Available at https://www.federalreserve.gov/econresdata/notes/feds-notes/2016/ and updated regularly.

[3] FRED series: DRTSCILM (C&I loans to large and middle-market firms) and DRTSCIS (C&I loans to small firms). Available at https://fred.stlouisfed.org/tags/series?t=sloos.

(2026).

Financial conditions indices. I use the Chicago Fed National Financial Conditions Index (NFCI) and its adjusted variant (ANFCI), which strips out the business-cycle component (FRED: NFCI, ANFCI). Both are weekly indices aggregated to monthly frequency. The ANFCI is available from 1971 and synthesizes 105 indicators of financial activity across money markets, debt markets, equity markets, and the banking sector. Positive values indicate tighter-than-average financial conditions. The ANFCI serves as my primary state variable for the state-dependent analysis because it isolates financial conditions from the business cycle—a distinction that is critical for disentangling the financial amplification channel from cyclical state dependence (Luna-Mallea, 2026). In robustness checks I present the state-dependent results using GZ credit spread as state variable.

3.3 Macroeconomic Variables

To replicate the Gavriilidis et al. (2026) baseline and to control for macroeconomic conditions, I include: industrial production (FRED: INDPRO), the unemployment rate (FRED: UNRATE), the Consumer Price Index for all urban consumers (FRED: CPIAUCSL), the 3-month Treasury bill rate (FRED: TB3MS), the effective federal funds rate (FRED: FEDFUNDS), the CBOE Volatility Index (FRED: VIXCLS), and the economic policy uncertainty index of Baker et al. (2016) from `policyuncertainty.com`. Industrial production and CPI enter in log levels ($100 \times \ln$), so that impulse responses capture cumulative percent changes. Additionally, I control for the Oil prices (WTI) to isolate the true effect of the CPU shocks, avoiding confounding with fluctuations in price for the carbon intensive sector.

3.4 Firm-Level Data

For the firm-level analysis, I construct a quarterly panel of U.S. publicly traded non-financial firms from Compustat North America, following the sample construction in Gavriilidis et al. (2026, Appendix C.2). I exclude financial firms (SIC 6000–6999) and utilities (SIC 4900–4999), and restrict the sample to firms with positive total assets and incorporated in the United States.

I construct the following variables from Compustat quarterly data: investment rate (capital expenditures divided by lagged net property, plant, and equipment: `capxy/L.ppentq`); log capital expenditures ($100 \times \ln(\texttt{capxy})$); R&D intensity (R&D expense over total assets: `xrdq/atq`); book leverage (total debt over total assets: $(\texttt{dlcq} + \texttt{dlttq})/\texttt{atq}$); cash-to-assets ratio (`cheq/atq`); return on assets (operating income before depreciation over total assets: `oibdpq/atq`); Tobin's Q proxy (market equity plus total debt divided by total assets); and log total assets as a measure of firm size. Observations with zero or missing capital expenditures are dropped from the log capex specifications; for the investment rate, observations with zero, negative, or missing lagged net PP&E are excluded. All ratios are winsorized at the 1st and 99th percentiles.

To study heterogeneity in climate change exposure, I merge the Sautner et al. (2023) firm-level climate change exposure data, available from OSF (https://osf.io/fd6jq/). The data provide quarterly measures of overall climate change exposure, as well as sub-components for opportunity, physical, and regulatory exposures, based on textual analysis of earnings call transcripts. The data cover over 10,000 firms worldwide from 2002 to 2020; I match to Compustat on `gvkey` and quarter.

I aggregate the monthly CPU shock to the quarterly frequency by summing within each calendar quarter to match the Compustat reporting frequency.

3.5 Summary Statistics

Table 1 presents summary statistics for the main variables used in the analysis. Panel A describes the CPU shock and the underlying index. The shock has zero mean and unit variance by construction. Panel B reports financial variables at a monthly frequency; the excess bond premium averages close to zero with substantial variation (standard deviation of 0.44 pp), while the GZ credit spread averages approximately 2.3 pp. Panel C describes the macroeconomic controls, and Panel D summarizes the firm-level Compustat variables.

4 Empirical Methodology

4.1 Aggregate Analysis: Local Projections

I estimate the dynamic causal effects of CPU shocks on financial and macroeconomic variables using local projections (Jordà, 2005). This approach is flexible, robust to misspecification of the data-generating process, and has become the standard method for estimating impulse responses in applied macroeconomics (Jordà, 2005; Stock and Watson, 2018; Plagborg-Møller and Wolf, 2021).

Table 1: Summary Statistics

Variable	Obs	Mean	Std. Dev.	Min	Median	Max
Panel A: Climate Policy Uncertainty						
CPU Shock (GKRS)	408	0.000	1.000	−5.31	−0.04	2.75
CPU Index (broad)	492	130.5	95.3	17.0	105.2	618.5
CPU Instrument	408	−0.01	0.69	−3.50	0.00	5.99
Panel B: Financial Variables						
Excess Bond Premium (pp)	408	−0.01	0.44	−0.98	−0.07	2.69
GZ Credit Spread (pp)	408	2.30	0.89	1.07	2.08	6.17
C&I Loans (log)	408	698.2	283.6	305.4	617.8	1,243
SLOOS Tightening (%)	119	3.8	19.4	−24.2	0.0	83.6
High Yield OAS (pp)	288	5.22	2.36	2.33	4.47	19.88
BBB OAS (pp)	288	1.83	0.75	0.96	1.62	6.01
ANFCI	408	−0.03	0.51	−0.72	−0.13	2.89
S&P 500 (log)	408	730.6	395.3	198.2	644.5	1,514
Panel C: Macroeconomic Variables						
Industrial Production (log)	408	448.2	11.3	416.8	449.1	469.4
Unemployment Rate (%)	408	5.72	1.53	3.40	5.50	10.00
CPI (Index, log)	408	510.4	20.8	465.2	513.6	549.1
3-Month T-Bill (%)	408	3.68	2.36	0.01	3.93	9.14
Panel D: Firm-Level Variables (Compustat quarterly)						
Investment Rate	196,428	0.072	0.078	0.001	0.048	0.498
Leverage	196,428	0.259	0.204	0.000	0.228	0.903
Cash/Assets	196,428	0.163	0.189	0.001	0.086	0.856
Climate Exposure	84,312	1.42	2.88	0.00	0.36	28.7

Notes: Sample period for aggregate analysis: 1986:M1–2019:M12 (408 monthly observations). SLOOS is quarterly (1990:Q2–2019:Q4, 119 quarters). Firm-level panel is quarterly from Compustat. CPU shock from Gavriilidis et al. (2026). EBP and GZ spread from the Federal Reserve Board. SLOOS and FRED series from the Federal Reserve Bank of St. Louis. Climate exposure from Sautner et al. (2023). C&I Loans and S&P 500 are in levels ($bn and index points, respectively); in the estimation they enter in $100 \times \ln$. All firm-level ratios are winsorized at the 1st and 99th percentiles.

4.1.1 Linear Local Projections

The baseline specification estimates the cumulative response of each outcome variable y at horizon h to a unit CPU shock:

$$y_{t+h} - y_{t-1} = \alpha_h + \theta_h \varepsilon_t^{\text{CPU}} + \sum_{\ell=1}^{p} \boldsymbol{\beta}_{h,\ell}' \mathbf{x}_{t-\ell} + \delta_{\text{year}(t)} + u_{t+h}, \quad h = 0, 1, \ldots, H, \tag{1}$$

where $\varepsilon_t^{\text{CPU}}$ is the Gavriilidis et al. (2026) structural CPU shock, $\mathbf{x}_{t-\ell}$ is a vector of lagged controls including the dependent variable and the CPU shock, and $\delta_{\text{year}(t)}$ denotes year fixed effects that absorb low-frequency trends and common annual shocks. The coefficient θ_h traces out the cumulative impulse response of y to a one-standard-deviation CPU shock at horizon h.

The left-hand side is the cumulative change $y_{t+h} - y_{t-1}$, so for variables in $100 \times \ln$ levels (industrial production, CPI, C&I loans, S&P 500), the IRF is directly interpretable as the cumulative percent change. For variables in levels (EBP, credit spreads, unemployment rate, interest rates), the IRF gives the cumulative change in percentage points.

To account for the dynamic transmission of the shock, I include lags of the dependent and control variables, with the optimal lag length selected via the Akaike Information Criterion (AIC). As demonstrated by Montiel Olea et al. (2025), this lag-augmentation sufficiently absorbs the serial correlation in the regression scores that is typically induced by the overlapping horizon structure of local projections. Consequently, valid inference does not require autocorrelation-consistent corrections; instead, standard errors are computed using a standard heteroskedasticity-robust estimator. I verify that the main results are robust to Newey-West standard errors with bandwidth $H + 1$ and to a fixed lag length of 12 months; the key impulse responses for EBP, GZ spread, and C&I loans retain significance under both alternatives. I report 90% and 68% confidence bands throughout the macro analysis; where results are significant only at the 68% level, I note this explicitly and interpret the evidence as suggestive rather than definitive.

The maximum horizon is set to $H = 36$ months for the aggregate analysis. I verify that results are robust to alternative horizons ($H = 24$, $H = 48$).

4.1.2 State-Dependent Local Projections

To test whether the financial effects of CPU depend on prevailing financial conditions, I estimate state-dependent local projections following the smooth-transition approach of Auerbach and Gorodnichenko (2012) and Tenreyro and Thwaites (2016):

$$y_{t+h} - y_{t-1} = \alpha_h + \theta_h^T F(z_{t-1})\, \varepsilon_t^{\text{CPU}} + \theta_h^L \left(1 - F(z_{t-1})\right) \varepsilon_t^{\text{CPU}} + \sum_{\ell=1}^{p} \boldsymbol{\beta}'_{h,\ell} \mathbf{x}_{t-\ell} + \delta_{\text{year}(t)} + u_{t+h}, \tag{2}$$

where $F(z_{t-1})$ is a smooth-transition function of the (lagged) state variable z_{t-1}:

$$F(z_{t-1}) = \frac{\exp(\gamma\, z_{t-1})}{1 + \exp(\gamma\, z_{t-1})}, \quad \gamma > 0. \tag{3}$$

The state variable z_{t-1} is the standardized (zero mean, unit variance) lagged ANFCI. Since positive ANFCI values indicate tighter-than-average financial conditions, $F(z) \to 1$ when financial conditions are tight ($z \gg 0$) and $F(z) \to 0$ when conditions are loose ($z \ll 0$).[4] The coefficients θ_h^T and θ_h^L capture the IRF in tight and loose financial conditions, respectively.

As a transparent baseline, I also report results from a simple binary interaction:

$$y_{t+h} - y_{t-1} = \alpha_h + \theta_h^T\, \mathbb{1}(z_{t-1} > 0)\, \varepsilon_t^{\text{CPU}} + \theta_h^L\, \mathbb{1}(z_{t-1} \leq 0)\, \varepsilon_t^{\text{CPU}} + \text{controls} + u_{t+h}, \tag{4}$$

where $\mathbb{1}(z_{t-1} > 0)$ indicates that the ANFCI is above zero (*i.e.*, financial conditions are tighter than average).

To formally test for state dependence, I evaluate the null hypothesis that the impulse responses are identical across regimes ($H_0 : \theta_h^T = \theta_h^L$). Rather than estimating standard period-by-period local projections—which complicates joint inference across horizons—I estimate a stacked local projection system (Jordà, 2005, 2009). By stacking the data and computing a single regression for the entire projection path, I recover the full variance-covariance matrix across all horizons. This allows me to perform a direct joint Wald test for the equality of the impulse responses across the entire horizon simultaneously, yielding a single, exact joint p-value without the need for multiple hypothesis testing corrections.

My baseline state variable is the ANFCI because it isolates financial conditions from business-cycle dynamics—a critical distinction given that Luna-Mallea (2026) documents state dependence on the business cycle. By using ANFCI, I can distinguish whether CPU's amplification operates through financial conditions *per se*, or merely proxies for cyclical state dependence. A potential concern is that one-period-lagged ANFCI may itself be correlated with anticipated climate-policy developments or with other omitted shocks that also shape the response path, since the ANFCI is a broad index built partly from the same credit-market conditions that serve as outcomes. To mitigate this concern, I verify that the main state-dependent results are robust to us-

[4]The smoothing parameter $\gamma = 1.5$ is set at the standard value in the literature. I show robustness to $\gamma \in \{1, 2, 3\}$.

ing a 4-quarter moving average of lagged ANFCI (which is more predetermined with respect to the shock) and to excluding observations during the Global Financial Crisis (2007:Q3–2009:Q2), which rules out the possibility that the tight-state results are driven entirely by one extreme stress episode. I present additional robustness using the EBP and GZ spread as alternative state variables in Section 5.

4.1.3 Identification

The identification of the causal effect of CPU rests entirely on the Gavriilidis et al. (2026) narrative instrument. The structural CPU shock $\varepsilon_t^{\text{CPU}}$ is extracted from a SVAR-IV in which the narrative event-based instrument satisfies two conditions:

1. *Relevance*: The instrument is correlated with the structural CPU shock. The first-stage F-statistic is approximately 23.5, well above the Stock and Yogo (2005) threshold of 10.
2. *Exogeneity*: The instrument is uncorrelated with other structural shocks in the system. Gavriilidis et al. (2026) demonstrate this in their Table D.2, showing that the narrative instrument is largely uncorrelated with financial market shocks, oil supply shocks, and monetary policy shocks.

By using the extracted structural shock as a direct regressor in my local projections, I inherit the GKRS identification. This "direct LP" approach is valid under the assumption that the SVAR correctly recovers the structural shock (Plagborg-Møller and Wolf, 2021).

Extension to financial outcomes. An important consideration is whether the GKRS identification extends to my financial outcome variables. The key requirement is that the narrative instrument is uncorrelated with shocks to the financial variables I study (EBP, credit spreads, bank lending). Since the GKRS events are driven by political decisions about climate policy—legislative actions, executive orders, judicial rulings, international agreements—rather than by financial market developments, exogeneity with respect to credit market shocks is plausible. To strengthen this claim, all baseline specifications control for lags of the VIX (which absorb common financial volatility shocks), lags of WTI oil prices (which account for energy-price confounders), and lags of the Baker et al. (2016) Economic Policy Uncertainty (EPU) index (which absorb broader policy uncertainty movements). With these controls, the remaining variation in the CPU shock is more clearly attributable to climate-specific policy uncertainty rather than to general financial, energy-market, or political conditions.

Distinctiveness from general uncertainty shocks. A natural concern is that the estimated credit market responses may simply reflect the generic effects of any uncertainty shock, rather than a channel specific to climate policy uncertainty. Existing evidence provides reassurance on this point. Luna-Mallea (2026) directly compares the macroeconomic dynamics under CPU shocks with those under general uncertainty measures—the VIX and the broad Equity Market Volatility (EMV) index of Baker et al. (2026)—and documents qualitatively different impulse response profiles. Whereas VIX and EMV shocks generate conventional demand-shock dynamics (output and prices fall together), CPU shocks produce a distinctive supply-shock pattern in which output falls while commodity prices and headline inflation rise simultaneously. This divergence in the joint price-output response cannot be generated by a common uncertainty channel and indicates that the GKRS narrative instrument captures a structurally distinct source of macroeconomic fluctuations. Moreover, the baseline specifications in this paper control for lagged VIX and EPU directly, so the estimated credit market responses are conditional on—and thus incremental to—general financial and policy uncertainty. The fact that CPU retains significant predictive power for credit spreads, bank lending, and the excess bond premium after absorbing these broad uncertainty controls supports the interpretation that climate policy uncertainty operates through a channel that is not subsumed by generic uncertainty dynamics.

4.2 Quarterly SLOOS Analysis

The SLOOS data on bank lending standards are available at quarterly frequency. To accommodate this, I estimate quarterly local projections:

$$\text{SLOOS}_{t+h} = \alpha_h + \theta_h\, \varepsilon_t^{\text{CPU},Q} + \sum_{\ell=1}^{4} \beta_{h,\ell}\, \text{SLOOS}_{t-\ell} + \sum_{\ell=1}^{4} \delta_{h,\ell}\, \varepsilon_{t-\ell}^{\text{CPU},Q} + \delta_{\text{year}(t)} + u_{t+h}, \quad h = 0, 1, \ldots, 12, \tag{5}$$

where $\varepsilon_t^{\text{CPU},Q}$ is the quarterly CPU shock, obtained by summing the monthly GKRS shock within each calendar quarter and re-standardizing to unit variance, so that impulse responses are comparable to a one-standard-deviation quarterly shock. The dependent variable is the net percentage of banks reporting tightening of C&I lending standards. The maximum horizon is 12 quarters.

4.3 Firm-Level Analysis

4.3.1 Average Firm-Level Response

I estimate the average firm-level investment response to CPU shocks using a panel local projection with firm fixed effects (Jordà, 2023):

$$y_{j,t+h} = \alpha_j + \theta_h\, \varepsilon_t^{\text{CPU},Q} + \sum_{\ell=1}^{p_h} \boldsymbol{\beta}'_{h,\ell} \mathbf{x}_{j,t-\ell} + u_{j,t+h}, \tag{6}$$

where $y_{j,t}$ is the log of investment (or R&D) for firm j in quarter t, α_j are firm fixed effects, and $\mathbf{x}_{j,t-\ell}$ includes lagged firm-level controls (log investment, Tobin's Q, cash flow, size) as well as lags of the shock. Following the econometric framework established by Almuzara and Sancibrián (2024) for estimating micro responses to macro shocks, I employ a time-clustered lag-augmented heteroskedasticity-robust (t-LAHR) inference strategy. Standard errors are clustered exclusively at the time (quarter) level, which fully accounts for the cross-sectional dependence induced by the aggregate nature of the CPU shock. Simultaneously, the inclusion of lagged controls ($\ell = 1 \ldots p_h$) absorbs the serial correlation in the regression scores, avoiding the need for unit-level clustering or autocorrelation-robust (HAR) corrections. To balance the need for generous lag augmentation with finite-sample degrees of freedom, the number of lags is selected dynamically for each projection horizon following the heuristic $p_h = \min\{h, \lfloor (T-h)^{1/3} \rfloor\}$, where T is the total time dimension of the sample (Almuzara and Sancibrián, 2024). As demonstrated by the authors, this combination of rule-based lag augmentation and time-clustering guarantees uniformly valid inference across horizons, regardless of the signal-to-noise ratio of the aggregate macroeconomic shock relative to the idiosyncratic firm-level variance.

4.3.2 Financial Heterogeneity: Leverage Interaction

The central firm-level test examines whether financially constrained or more climate-exposed (Sautner et al., 2023) firms are disproportionately affected by CPU shocks. Following Ottonello and Winberry (2020) and the heterogeneity specification in Gavriilidis et al. (2026), I estimate:

$$y_{j,t+h} - y_{j,t-1} = \alpha_j + \delta_t + \gamma_h \left(\mathrm{X}_{j,t-1} - \overline{\mathrm{X}}_j \right) \times \varepsilon_t^{\mathrm{CPU},Q} + \boldsymbol{\beta}_h' \mathbf{x}_{j,t-1} + u_{j,t+h}, \tag{7}$$

where $\mathrm{X}_{j,t-1}$ is a lagged firm characteristic—either book leverage or climate change exposure—and $\overline{\mathrm{X}}_j$ is the corresponding firm-level time-series mean of X, so that the interaction exploits *within-firm* variation in the characteristic over time. Both firm (α_j) and time (δ_t) fixed effects are included; the time fixed effects absorb the aggregate CPU shock and any common confounders, so the coefficient γ_h captures the *differential* effect of CPU on firms with temporarily above-average X. The demeaned characteristic is standardized to unit variance for ease of interpretation: γ_h gives the additional investment response for a one-standard-deviation increase in within-firm X.

I also examine interactions using a triple interaction between climate exposure and leverage to assess whether the effects are amplified for firms that are both more exposed to climate transition risk and more financially constrained. This specification allows me to test whether leverage strengthens the transmission of climate policy uncertainty by tightening credit conditions for firms whose assets are more vulnerable to changes in climate policy. If the financial channel is operative, the contractionary effects of CPU shocks should be strongest among highly leveraged firms with greater exposure to climate risk, consistent with a mechanism in which declines in the collateral value of brown assets amplify borrowing constraints.

Standard errors are clustered at the quarter level throughout the firm-level analysis, following Almuzara and Sancibrián (2024)

5 Aggregate Results

This section presents the main empirical findings in four parts. Section 5.1 documents the aggregate financial responses to CPU shocks at monthly frequency, establishing the sequencing of financial transmission. Section 5.2 corroborates these findings at quarterly frequency and introduces the bank lending standards channel via the SLOOS. Section 5.3 presents mediation analysis testing whether observed lending standards explain the credit contraction. Section 5.4 examines state dependence on financial conditions.

5.1 Aggregate Financial Responses

Figure 1 reports impulse responses to a one-standard-deviation CPU shock estimated via local projections at monthly frequency over a 36-month horizon. Shaded areas denote 68% and 90% confidence bands computed with robust standard errors.

Real activity and the policy rate. Investment (private fixed investment) declines on impact and falls steadily over the first 12–16 months, reaching a trough of approximately -0.5% around months 16–20. The response remains negative and statistically significant at the 68% level through much of the horizon, before partially reverting beyond month 24. The effective federal funds rate declines promptly, falling by roughly 10–15 basis points within the first 6 months and remaining persistently negative through the full 36-month horizon. This pattern is consistent with an endogenous monetary policy easing in response to the contractionary effects of the CPU shock, as also documented in the quarterly analysis of Gavriilidis et al. (2026).

Credit quantities. C&I loans at commercial banks exhibit a gradual and persistent contraction. The decline begins around months 4–6, reaches approximately -0.5% by month 12, and deepens to roughly -0.8 to -1.0% by months 20–24. The response is statistically significant at the 90% level from approximately month 8 onward and shows no meaningful reversion over the 36-month horizon. The magnitude is economically substantial: in a banking system with roughly $2.3 trillion in C&I loans, a 1% contraction represents approximately $23 billion in reduced credit availability.

Credit pricing. The GZ credit spread—which captures both default risk and risk premia—increases immediately following the shock, reaching approximately 5–8 basis points by months 12–18, with the response significant at the 68% level throughout most of the first two years. The excess bond premium (EBP) rises by a similar magnitude, peaking at roughly 5–8 basis points around months 4–12 before gradually reverting toward zero. The expected default component displays a smaller positive response on the order of 2 basis points, but the estimate is imprecise and statistically insignificant at conventional levels for most of the horizon. The contrast between the EBP and expected default responses suggests that the initial widening of credit spreads is driven primarily by an increase in risk premia—reflecting the risk-bearing capacity of financial intermediaries—rather than by a deterioration in borrower fundamentals, consistent with the interpretation of the EBP as a financial friction indicator (Gilchrist and Zakrajšek, 2012).

Corporate bond spreads and financial conditions. Across the risk spectrum, corporate credit spreads widen persistently. The ICE BofA high-yield option-adjusted spread increases by approximately 15–25 basis points over months 12–20, representing a meaningful repricing of credit risk for speculative-grade borrowers. The investment-grade corporate OAS (broad U.S.) rises by about 5–10 basis points over the same horizon, confirming that the spread widening extends beyond high-yield into the broader credit market. The Chicago Fed NFCI tightens gradually, rising by approximately 3–4 basis points through months 12–20, corroborating the picture of a broad tightening in financial conditions.

Relative timing. A notable feature of the monthly results is the relative timing of financial responses. Risk premia respond earliest: the EBP and GZ credit spread rise within the first few months, peaking around months 4–12. Corporate credit spreads follow with a slight delay, widening most sharply between months 8 and 20. Credit quantities adjust later: C&I loans begin declining around months 4–6 and continue deepening through months 20–24, well after risk premia have begun to revert. This pattern of relative timing—pricing moving before quantities—is *consistent with* a transmission mechanism in which CPU shocks initially raise intermediary risk premia and widen corporate spreads, which in turn lead banks to tighten credit supply and reduce loan volumes. However, with overlapping LP horizons and different degrees of persistence across series, these onset differences should be interpreted as suggestive co-movement patterns rather than as formally identified causal sequencing. The monetary policy response (declining federal funds rate) is roughly contemporaneous with the rise in risk premia, suggesting that the central bank reacts to the emergent tightening in financial conditions but cannot fully offset the credit contraction.

Figure 1: Aggregate Financial Responses to a CPU Shock — Monthly

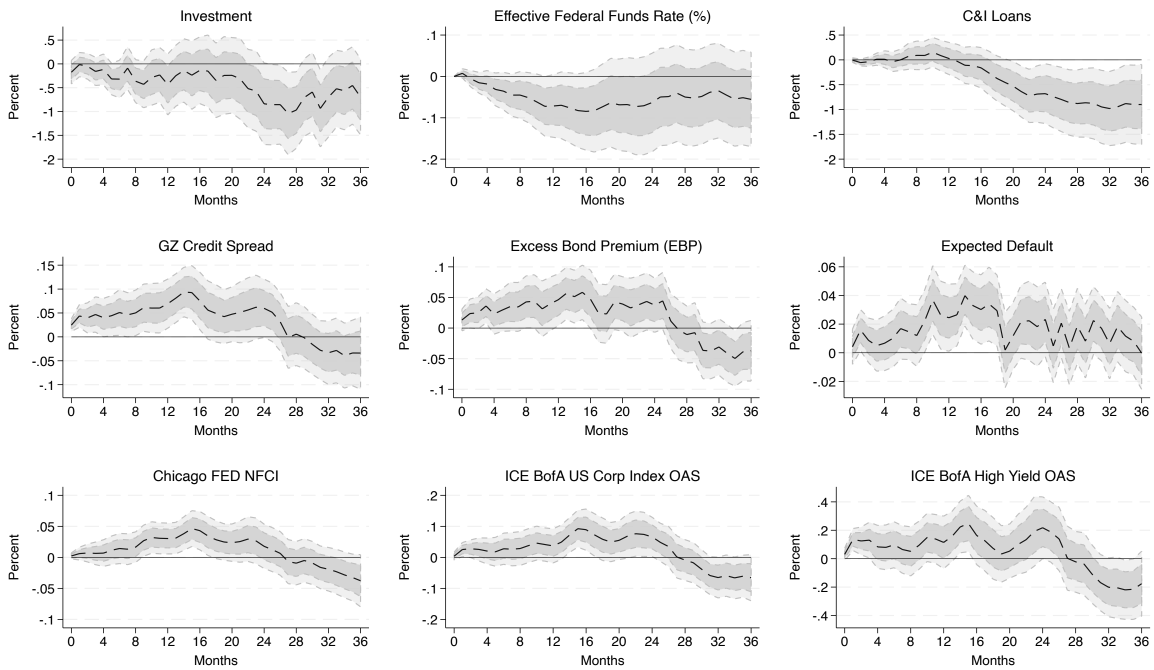

Notes: Impulse responses to a one-standard-deviation climate policy uncertainty shock, estimated via local projections (equation 1). The shock is the structural CPU shock from Gavriilidis et al. (2026). Controls include 12 lags of the dependent variable and the CPU shock, plus year fixed effects. Shaded areas denote 68% (dark) and 90% (light) confidence bands.

Robustness to alternative CPU index. Additionally, the results are robust to using the Luna-Mallea (2026) PCA-based CPU index, which combines the Environmental Policy Uncertainty index of Palikhe et al. (2024) and the Energy and Environmental Regulation subindex from the Equity Market Volatility Tracker (Baker et al., 2026) via principal component analysis. This index captures a broader set of news-based signals about climate policy uncertainty than the GKRS measure alone, though it lacks the narrative-instrument-based identification. When I replace the GKRS structural shock with the AR(1) innovation of the first principal component of the Luna-Mallea (2026) CPU index after c in the baseline local projections, the main findings are preserved,

5.2 Quarterly Evidence and the Bank Lending Channel

To incorporate the Senior Loan Officer Opinion Survey—available only at quarterly frequency—and to facilitate comparison with the firm-level analysis, I re-estimate the local projections at quarterly frequency over a 12-quarter horizon. Figure 2 reports the results.

Real activity and credit quantities. The quarterly estimates corroborate the monthly findings. Private investment declines gradually, reaching approximately −0.8 to −1.0% around quarters 6–8, and remains negative through the full horizon. The effect is statistically significant at the 68% confidence level for most of the horizon. The effective federal funds rate shows a persistent decline of roughly 20–25 basis points, beginning almost immediately and lasting through the full 12-quarter horizon, though significance is concentrated at the 68% band. C&I loans exhibit a persistent contraction reaching approximately −2.0% by quarters 8–10, highly significant at the 90% level throughout—confirming the monthly baseline at a different frequency.

Credit pricing. The GZ credit spread displays a positive response in the first two years, peaking at approximately 8 basis points around quarters 3–5, significant at the 68% level. The excess bond premium increases by roughly 8 basis points, with significance concentrated in the first four quarters. The expected default component is economically small and statistically insignificant across the full horizon at quarterly frequency. Taken together, the quarterly pricing evidence suggests that the initial credit spread widening documented in the monthly analysis is driven primarily by the risk-premia component (EBP) rather than by a deterioration in expected fundamentals.

Bank lending standards (SLOOS). The SLOOS results provide direct evidence on the credit supply mechanism. I examine three dimensions: loan demand reported by banks, and lending standards for medium-to-large firms and for small firms.

The net percentage of banks tightening standards for medium and large firms rises sharply in the first one to three quarters, peaking at approximately 2–3 percentage points, before fading toward zero and eventually turning negative (net easing) around quarters 8–10. Tightening for small firms follows a similar pattern but is more persistent and larger in magnitude, peaking at roughly 3 percentage points around quarter 4. The stronger and more persistent tightening for small firms is consistent with a financial accelerator mechanism: smaller borrowers with weaker collateral and fewer outside options face more severe credit rationing when banks become cautious about climate policy risk.

Loan demand reported by banks drifts negative, reaching approximately −2 to −3 percentage points through quarters 4–8, before partially recovering. However, the demand response is noisier and largely statistically insignificant at the 90% level. The contrast is informative: the credit *supply* response (standards tightening) is sharper, earlier, and more precisely estimated than the *demand* response, consistent with the initial credit contraction being supply-driven.

Supply-side interpretation. The SLOOS evidence supports a credit supply interpretation of the C&I loan decline. Standards tighten early and sharply (quarters 1–4), preceding the bulk of the loan contraction (which deepens through quarters 6–10). The subsequent normalization of lending standards—and their eventual move to net easing around quarters 8–10—suggests that the initial tightening episode is transitory. Yet the C&I loan contraction persists well beyond the period of active tightening, indicating that the initial supply shock sets in motion a credit contraction that is sustained by additional forces, including reduced investment demand and general equilibrium feedback.

Figure 2: Quarterly Impulse Responses to a CPU Shock

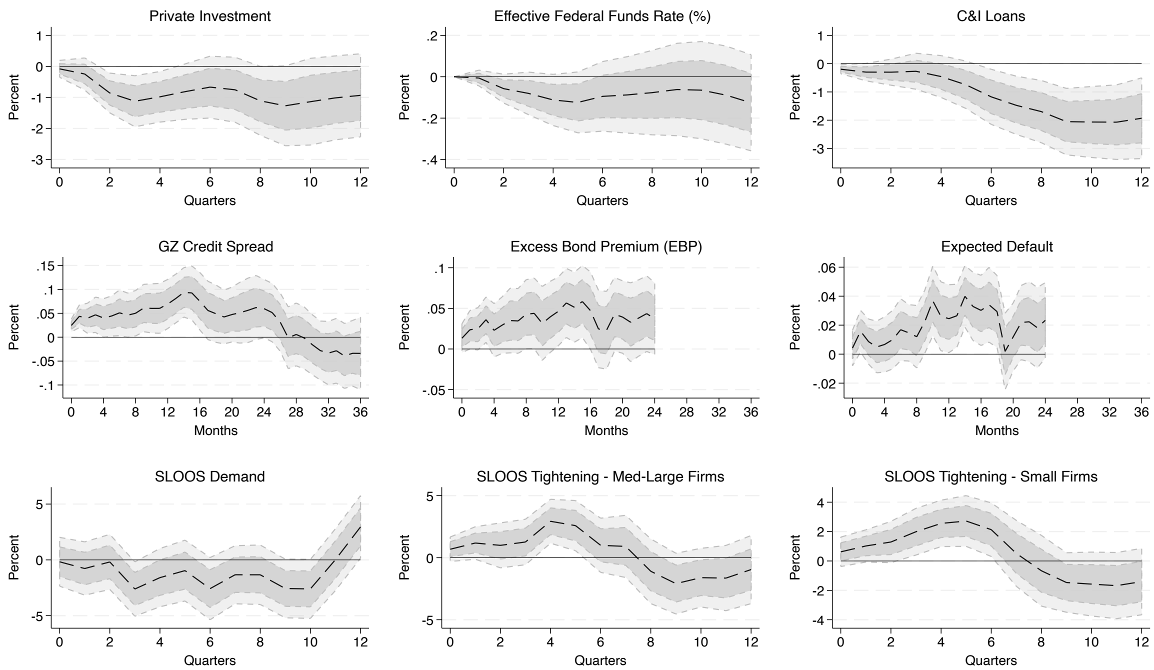

Notes: Impulse responses to a one-standard-deviation climate policy uncertainty shock at quarterly frequency. The quarterly CPU shock is the sum of monthly GKRS shocks within each calendar quarter. Controls include 4 lags of the dependent variable and the CPU shock, plus year fixed effects. Robust Standard Errors. Shaded areas denote 68% (dark) and 90% (light) confidence bands. SLOOS series begin in 1990:Q2.

5.3 Decomposition: Lending Standards and the Credit Contraction

The SLOOS evidence in the previous subsection establishes that CPU shocks tighten bank lending standards. A natural question is whether this tightening is associated with the subsequent decline in C&I loans. To assess this, I augment the baseline C&I loan specification with contemporaneous and lagged SLOOS tightening indicators for large and middle-market firms and re-estimate the local projections. Figure 3 reports the baseline (left panel) and SLOOS-controlled (right panel) responses. This exercise is a descriptive decomposition—SLOOS is itself an endogenous survey measure that captures both supply- and demand-side conditions—and should be interpreted as an attenuation exercise rather than a formal mediation design with identified causal shares.

Short-run attenuation. The results reveal a striking pattern. In the baseline specification, C&I loans decline monotonically from the onset, reaching a trough of approximately -2.0 to -2.1% at quarters 8–10. When SLOOS controls are added, the response in the first three quarters changes markedly: the point estimate turns slightly positive (peaking at roughly $+0.5\%$ around quarters 2–3), though it remains statistically insignificant at the 90% level. This indicates that, conditional on observed lending standards, the short-run credit contraction is absorbed—the variation in C&I loans that the CPU shock explains in the first few quarters is closely associated with the survey-measured tightening of bank lending standards.

Long-run persistence. In contrast, the medium-to-long-run response is largely unchanged by the inclusion of SLOOS controls. From quarter 4 onward, the SLOOS-controlled specification yields a C&I loan response that is qualitatively and quantitatively similar to the baseline, reaching a trough of approximately -2.0 to -2.2% around quarters 8–10. The confidence bands are wider, as expected given the additional controls, but the point estimates overlap closely.

Interpretation. The decomposition results are consistent with a two-phase interpretation of the credit channel. In the *short run* (quarters 0–3), the credit contraction co-moves closely with the tightening of observed lending standards: conditioning on SLOOS absorbs the short-run loan response. In the *medium run* (quarters 4–12), the loan contraction persists and deepens even after controlling for observed lending standards. This longer-run persistence likely reflects forces beyond the measured supply channel: reduced loan demand as firms scale back investment plans in response to both uncertainty and tighter credit, general equilibrium feedback from the output contraction, and potentially unmeasured dimensions of credit tightening (such as informal credit rationing, collateral haircuts, or relationship-based lending adjustments) that are not captured by the SLOOS survey.

This pattern—short-run attenuation and long-run persistence—is informative for the mechanisms discussion in Section 7. It suggests that bank lending standards are an important correlate of the *initial* credit contraction, but that the sustained decline reflects a broader set of channels, consistent with the financial accelerator dynamics modeled in Luna-Mallea (2026).

Figure 3: C&I Loan Response: Baseline vs. SLOOS-Controlled

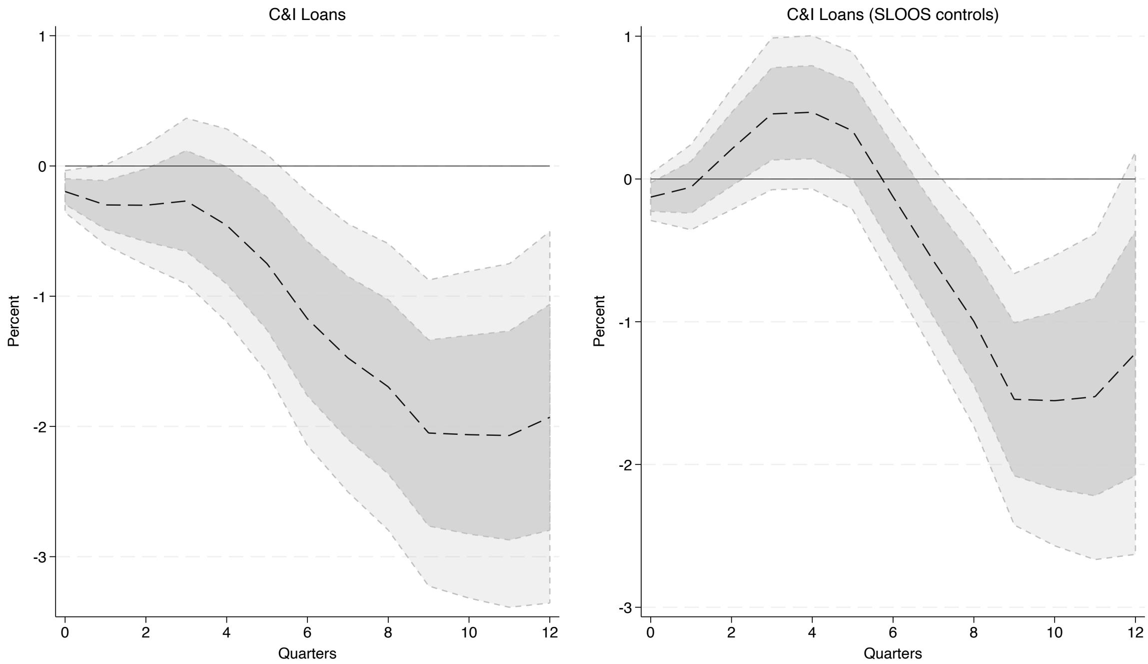

Notes: Left panel: baseline quarterly local projection of C&I loans on the CPU shock. Right panel: same specification augmented with contemporaneous and four lags of the SLOOS net tightening indicator for C&I loans to large and middle-market firms (FRED: DRTSCILM). Shaded areas denote 68% (dark) and 90% (light) confidence bands.

5.4 State Dependence on Financial Conditions

The results thus far indicate that CPU shocks, on average, transmit through credit markets. A natural extension is to examine whether these effects are amplified when the financial system is already under stress. I estimate the state-dependent local projections described in Section 4.1.2, using the Adjusted NFCI as the baseline state variable.

Figure 4 reports the state-dependent impulse responses, where the solid blue denotes the response during loose financial conditions (ANFCI ≤ 0) and the dashed red line denotes the response during tight conditions (ANFCI > 0). Below each panel, I report the p-value for the null hypothesis of equal responses across states, and the significance of each state from the F-test.

Real activity and monetary policy. The results show that the contractionary effects of a one standard deviation CPU shock are concentrated in periods of tight financial conditions. Investment falls more sharply and remains depressed for longer in the tight regime, while the policy response is also state dependent: the federal funds rate declines meaningfully only when credit conditions are already adverse, consistent with monetary accommodation aimed at cushioning the tightening in financial markets. Credit quantities exhibit an opposite pattern across regimes. In tight states, C&I loans decline significantly and persistently, whereas in loose states, credit is broadly stable and may even rise slightly at short horizons, with effects that are economically small and statistically weak. Taken together, these dynamics indicate that CPU shocks trigger a financial amplification mechanism that operates primarily when balance sheets are fragile and credit supply is more likely to tighten, translating a rise in uncertainty into a larger and more persistent contraction in investment and bank lending.

Credit pricing. State dependence in credit pricing is pronounced. The excess bond premium increases significantly in tight-credit states, by roughly 20 to 30 basis points, whereas the response under loose conditions is economically small and statistically indistinguishable from zero. A similar pattern holds for the GZ credit spread: it widens by about 10 to 20 basis points in tight states but remains essentially flat in loose states, and the null of equal responses across regimes cannot be rejected at conventional levels ($p(\text{equal}) = 0.12$). The expected-default component moves in the opposite direction, but the estimates are generally imprecise, particularly in the tight regime. Corporate bond spreads display even sharper asymmetries. High-yield option-adjusted spreads widen by approximately 80 to 100 basis points during tight conditions, with a much more muted reaction in loose states, and the broad U.S. corporate OAS follows the same state-dependent pattern, with equality strongly rejected ($p(\text{equal}) = 0.00$).

Bank lending standards and loan demand (SLOOS). The SLOOS responses provide direct evidence that the transmission of a CPU shock is strongly state dependent. In tight financial conditions (red dashed line), loan demand drops sharply on impact, reaching declines on the order of 10 percentage points within the first year, and the null of no effect is rejected ($p(\text{tight} = 0) = 0.00$). In contrast, demand is much more stable in loose conditions (blue line), with small movements around zero and no clear sustained contraction ($p(\text{loose} = 0) = 0.30$), implying a statistically meaningful asymmetry across regimes ($p(\text{equal}) = 0.00$). Lending standards display the same pattern but with clearer heterogeneity by borrower size. In tight states, banks report a pronounced tightening for both medium and large firms and, even more strongly, for small firms, with tightening peaking around 7–8 percentage points early in the horizon and remaining statistically significant for small firms ($p(\text{tight} = 0) = 0.00$). Under loose conditions, the tightening responses are comparatively muted and less precisely estimated, especially for medium and large firms ($p(\text{loose} = 0) = 0.52$), while equality is rejected for small firms ($p(\text{equal}) = 0.01$). Overall, the SLOOS evidence indicates that CPU shocks coincide with a sharp deterioration in credit demand and a marked tightening of credit supply in adverse financial states, with the supply response concentrated among small borrowers, consistent with a financial accelerator operating through bank intermediation.

Interpretation. The ANFCI-based state dependence reveals that the financial transmission of CPU documented in the preceding subsections is heavily concentrated in periods of tight financial conditions. When credit markets are already stressed, CPU shocks generate a powerful amplification: risk premia spike and credit spreads widen sharply (particularly for lower-rated borrowers) more than during normal conditions. This pattern is consistent with a financial accelerator mechanism in which CPU shocks interact with existing balance sheet vulnerabilities to produce nonlinear financial disruption—the "doom loop" between policy uncertainty and financial stress hypothesized in Section 1.

Importantly, this financial-conditions state dependence differs from the business-cycle state dependence documented in Luna-Mallea (2026), who finds larger CPU effects during expansions. The ANFCI strips out the business-cycle component, isolating financial conditions *per se*. The finding that financial conditions—rather than the output gap—drive the amplification provides a more precise test of the financial transmission mechanism proposed in the theoretical framework of Luna-Mallea (2026).

I present robustness to alternative state variables (excess bond premium and GZ credit spread) in Appendix A. The EBP-based results yield qualitatively similar but statistically weaker state dependence, while the GZ-based split produces a complementary pattern in which some variables respond more strongly during *loose* credit-spread states—consistent with a vulnerability interpretation in which financial disruption is larger when markets are complacent. Additionally, the results are robust to using the Luna-Mallea (2026) PCA CPU index.

6 Firm-Level Results

The aggregate results in Section 5 establish that CPU shocks trigger a persistent contraction in C&I lending and private investment, accompanied by a transitory rise in risk premia—the GZ credit spread increases by approximately 8 basis points and the excess bond premium by roughly 8 basis points in the first year, before reverting toward zero. A natural follow-up is whether these aggregate dynamics mask important heterogeneity across firms. In this section, I exploit firm-level variation in climate change exposure, sector composition, and financial leverage to trace the transmission of CPU shocks to corporate capital expenditure and R&D spending.

Figure 4: State-Dependent Financial Responses: Tight (red) vs. Loose (blue) Financial Conditions (ANFCI)

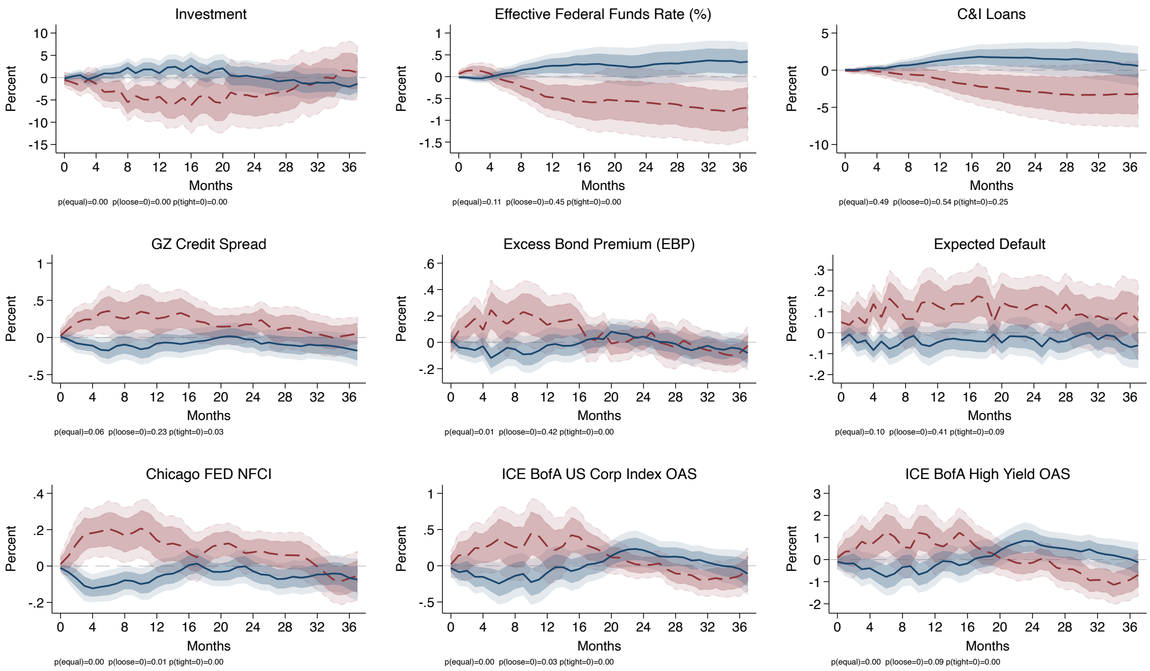

Notes: State-dependent impulse responses estimated via equation (2). The p-value for equality of responses across states is computed at each horizon via a Wald test; the reported $p(\text{equal})$ is the joint test across all horizons. Robust standard errors. Shaded areas denote 68% and 90% confidence bands.

Figure 5: State-Dependent Financial Responses: Tight vs. Loose Financial Conditions (ANFCI) — SLOOS

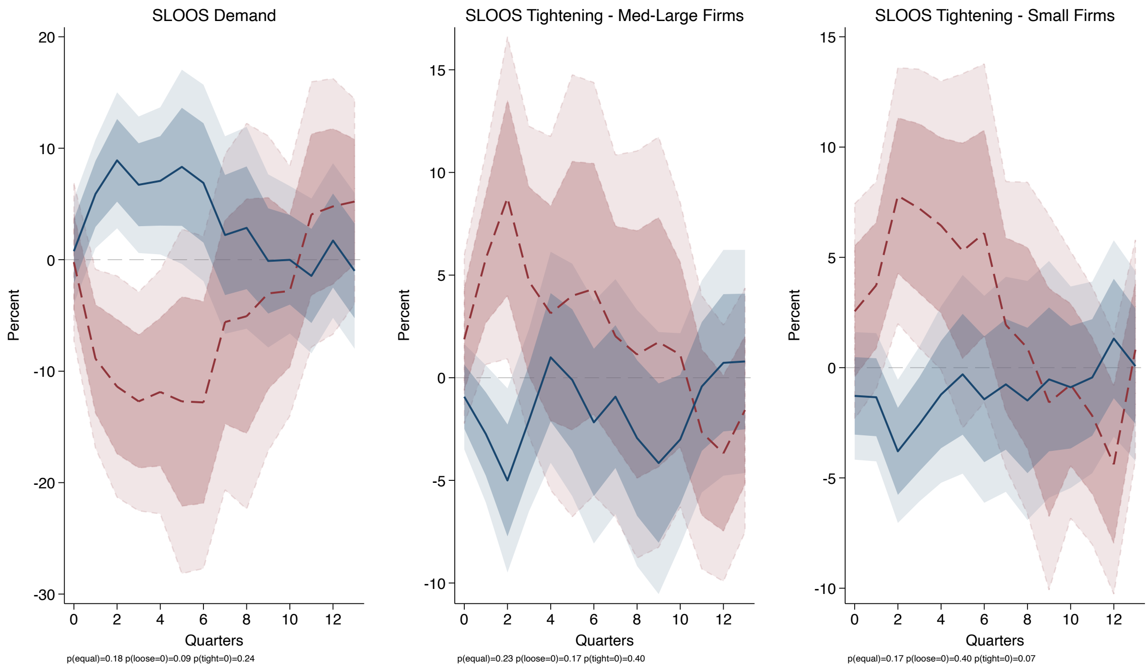

Notes: State-dependent impulse responses estimated via equation (2) at quarterly frequency, using the SLOOS lending standards and loan demand measures. The p-value for equality of responses across states is computed at each horizon via a Wald test; the reported $p(\text{equal})$ is the joint test across all horizons. Robust standard errors. Shaded areas denote 68% and 90% confidence bands.

The firm-level analysis uses two specifications introduced in Section 4.3. For baseline average effects (Section 6.1), I estimate equation (6) without time fixed effects, so the aggregate CPU shock directly enters as a regressor, with standard errors clustered at the date level following Almuzara and Sancibrián (2024). For heterogeneity analysis (Sections 6.2–6.4), I estimate equation (7) with both firm and time fixed effects, where $X_{i,t-1}$ denotes a lagged firm characteristic (climate change exposure, leverage, or their interactions). Time fixed effects absorb the aggregate CPU shock and any common confounders, so the interaction coefficient γ_h captures the *differential* effect of CPU on firms with temporarily above-average X relative to the average firm. Standard errors in the heterogeneity specifications are two-way clustered at the firm and date level to account for both cross-sectional and temporal dependence.

6.1 Baseline Firm-Level Responses

Before turning to heterogeneity, I document the average firm-level response to CPU shocks. Figure 6 presents impulse responses from a specification without time fixed effects, allowing the aggregate CPU shock to directly affect firm outcomes.

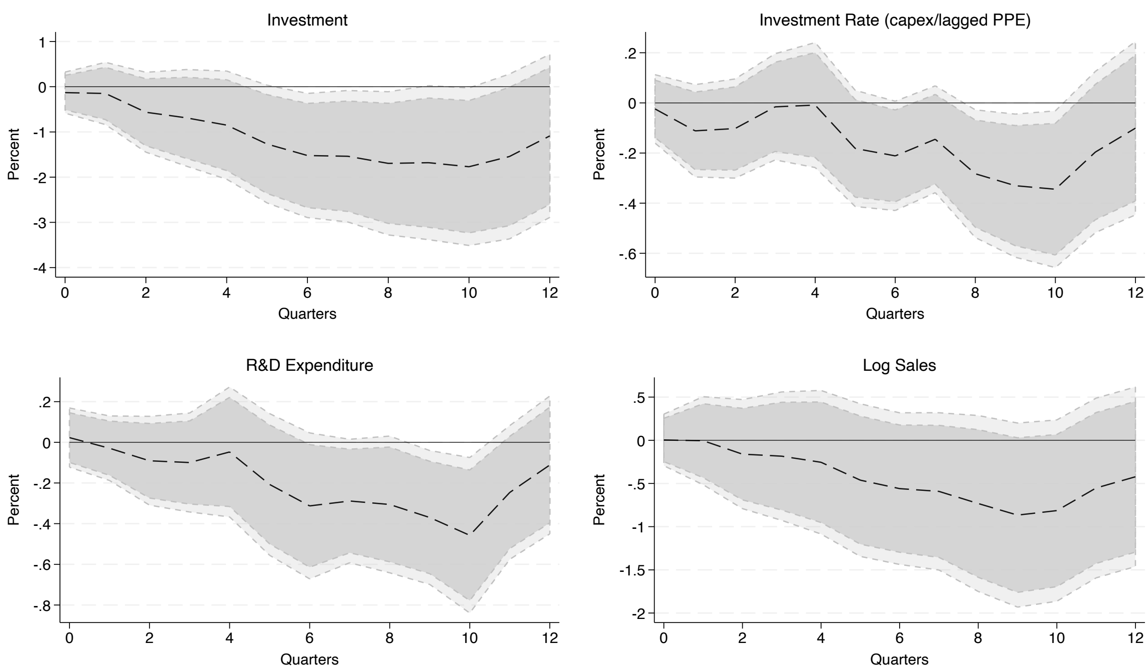

Figure 6: Baseline Firm-Level Responses to CPU Shocks

Notes: Impulse responses from firm-level local projections as in equation (6) without time fixed effects. Top left: log capital expenditure. Top right: investment rate (capex/lagged PPE). Bottom left: log R&D expenditure. Bottom right: log sales. Shaded areas denote 90% and 95% confidence bands based on standard errors clustered at the date level.

The results confirm that the aggregate-level findings in Section 5 are reflected in firm-level data. Log capital expenditure declines steadily following the CPU shock, reaching approximately -1.5% around quarters 8–10 before partially reverting by quarter 12. The investment rate (capex/lagged PPE) follows a similar trajectory, falling by roughly 0.3 percentage points around quarter 5, with significance at the 90% level during the middle of the horizon. R&D expenditure also contracts, though with a lag relative to physical investment: log R&D declines gradually, reaching approximately -0.4% around quarters 8–10, significant at the 90% level from approximately quarter 4 onward. Log sales decline by roughly 0.8% by quarters 8–10, consistent with the aggregate demand contraction documented at the macro level. The relative magnitudes indicate that firms cut physical investment more sharply and earlier than R&D spending, consistent with a "wait-and-see" interpretation in which firms prioritize protecting long-run innovation capacity while scaling back near-term capital formation (Bloom, 2009).

6.2 Climate Exposure and Financial Leverage

I now turn to the central heterogeneity analysis. Figure 7 reports interactions between the CPU shock and climate change exposure from Sautner et al. (2023). The top row shows the specification with firm and time fixed effects, where the coefficient captures purely cross-sectional heterogeneity: how firms with above-average climate change exposure respond differentially to CPU shocks relative to the average firm.

Climate change exposure. The interaction of CPU shocks with climate change exposure is negative for investment, with the point estimate reaching approximately -0.5 to -0.8% by quarters 6–10 (Figure 7, top-left). This indicates that firms with one-standard-deviation higher climate change exposure experience an additional investment decline of roughly half a percentage point on top of the average effect. For R&D expenditure (top-right), the interaction is negative in the early horizons but becomes imprecise and turns slightly positive beyond quarter 6, suggesting that climate-exposed firms may redirect resources toward adaptive or green R&D at longer horizons. The middle row (no time fixed effects) confirms this general pattern, while the bottom row shows the direct aggregate effect of CPU shocks on investment and R&D in the absence of time fixed effects: investment falls by approximately 2–3% over 10–12 quarters and R&D declines by about 0.5%, both significant at the 90% level.

Book leverage. Figure 8 presents the corresponding analysis for book leverage. The interaction between the CPU shock and within-firm demeaned leverage is *positive* for investment (top-left panel), with the coefficient rising from approximately 0.1% on impact to roughly 0.3–0.4% at quarters 6–10. This positive sign indicates that firms with temporarily above-average leverage cut investment *less* than firms with below-average leverage following CPU shocks. For R&D (top-right), the leverage interaction is also mildly positive but economically small and statistically imprecise.

This finding may appear counterintuitive—one might expect that highly leveraged firms, being more financially constrained, would reduce investment more. However, it is directly consistent with Ottonello and Winberry (2020), who show that low-leverage firms are more responsive to aggregate shocks because they operate farther from their borrowing constraint and thus have greater financial flexibility to adjust. Under this interpretation, when CPU shocks hit, less-constrained firms optimally scale back investment as a "wait-and-see" response (Bloom, 2009), while highly leveraged firms—already near their constraint boundary—face limited room for further adjustment. The middle row (no time FE) corroborates this pattern, and the bottom row confirms a large negative direct effect of CPU shocks on both investment and R&D in the specification without time fixed effects.

Triple interactions: climate exposure, leverage, and CPU. Figure 9 reports specifications that jointly interact the CPU shock with climate change exposure and leverage. The top panel (CC Exposure $\times$ CPU) confirms that climate-exposed firms experience differentially larger investment declines, with a negative coefficient reaching approximately -0.5 to -1.0% by quarters 4–8. The middle panel (Financial Constraint $\times$ CPU) is imprecise, oscillating around zero without sustained significance. The triple interaction in the bottom panel (CPU $\times$ Leverage $\times$ CC Exposure) is also statistically imprecise over most of the horizon, though it trends negative at longer horizons, providing suggestive—but inconclusive—evidence that the combination of high climate exposure and high leverage may amplify investment declines at extended horizons.

6.3 Summary of Firm-Level Evidence

The firm-level results can be summarized in three main findings. First, CPU shocks generate a broad-based contraction in firm investment, R&D, and sales, with magnitudes that are quantitatively consistent with the aggregate evidence from Section 5. Capital expenditure falls by approximately 1.5% and R&D by roughly 0.4% over a 2–3 year horizon. Second, firms with greater climate change exposure—as measured by the Sautner et al. (2023) textual indicators—experience differentially larger investment declines, confirming that CPU transmits more forcefully to firms with higher transition risk. Third, the leverage interaction reveals a nuanced pattern consistent with Ottonello and Winberry (2020): less-leveraged firms cut investment more sharply following CPU shocks, reflecting their greater financial flexibility to adjust, while highly leveraged firms are more constrained in their ability to change investment. Importantly, the bond market evidence (Section 6.4) shows that leverage *does* amplify the

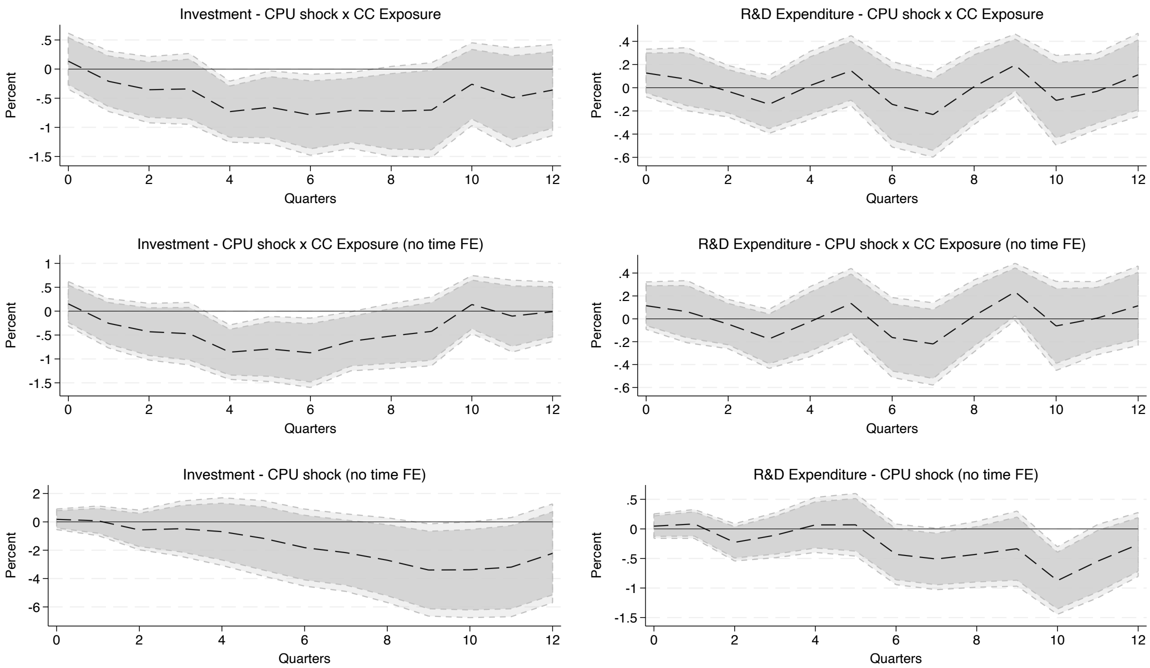

Figure 7: Firm-Level Responses to CPU Shocks: Climate Exposure
Notes: Impulse responses from firm-level local projections as in equation (7). Top row: interaction with relative climate change exposure. Middle row: interaction with a brown-sector indicator. Bottom row: triple interaction. Left column: log capital expenditure. Right column: log R&D expenditure. Shaded areas denote 90% and 95% confidence bands based on standard errors clustered at the date level.

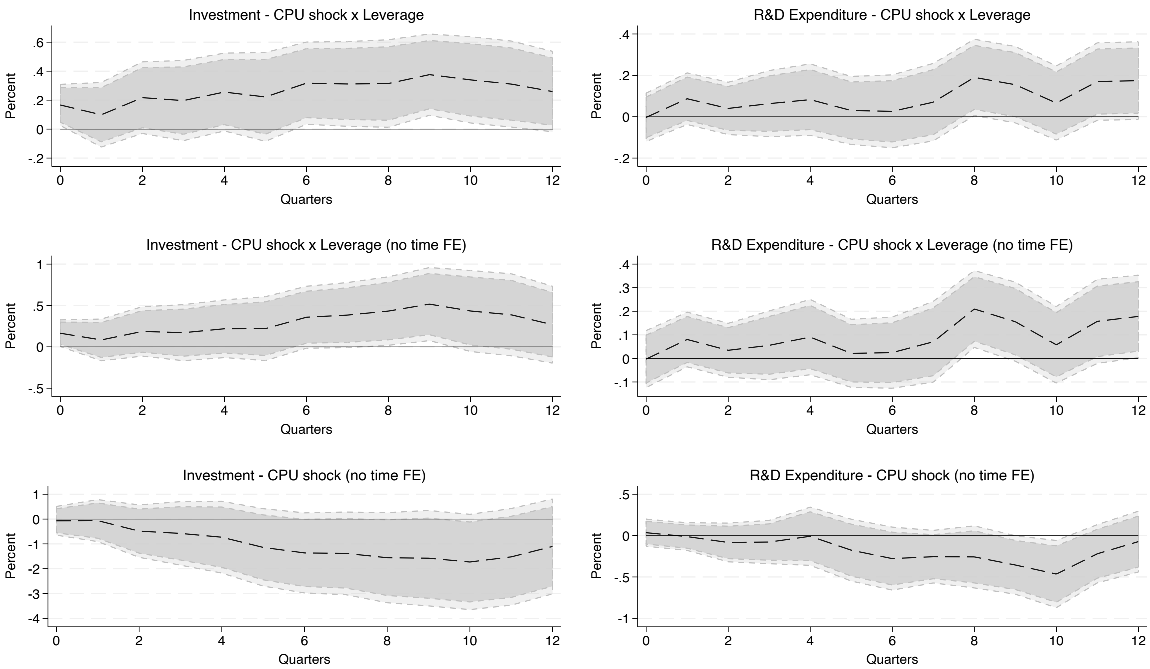

Figure 8: Firm-Level Responses to CPU Shocks: Book Leverage
Notes: Impulse responses from firm-level local projections as in equation (7). Top row: interaction with relative climate change exposure. Middle row: interaction with a brown-sector indicator. Bottom row: triple interaction. Left column: log capital expenditure. Right column: log R&D expenditure. Shaded areas denote 90% and 95% confidence bands based on standard errors clustered at the date level.

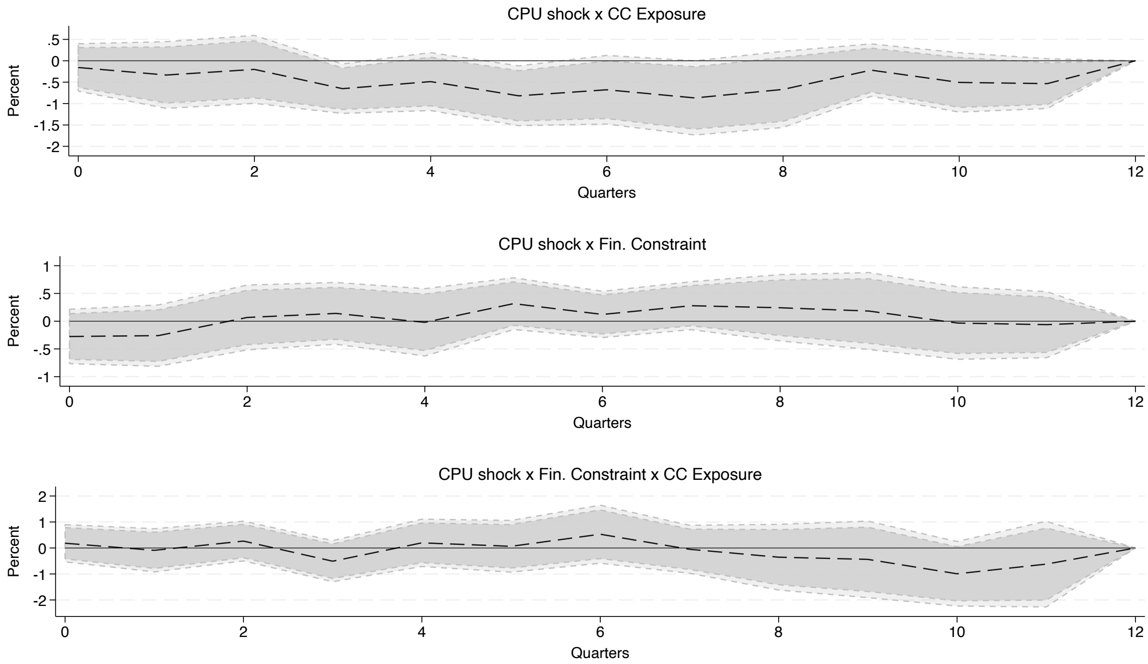

Figure 9: Firm-Level Responses: Climate Exposure and Leverage Interactions

Notes: Top row: leverage interaction. Middle row: brown-sector interaction. Bottom row: triple interaction. Left column: log capital expenditure. Right column: log R&D expenditure. Shaded areas denote 90% and 95% confidence bands based on standard errors clustered at the date level.

financial penalty—higher-leverage firms face rising bond yields and declining bond volumes—even though the investment response is more muted. This suggests that the financial accelerator operates primarily through the cost and availability of credit rather than through real investment adjustment for the most constrained firms.

6.4 Bond Market Evidence

The aggregate results in Section 5.1 document a widening of corporate credit spreads (GZ spread, high-yield OAS) and a rise in the excess bond premium following CPU shocks. To trace these aggregate credit market dynamics to the firm level, I exploit individual corporate bond data from WRDS/Mergent FISD, constructing firm-level measures of bond yields (value-weighted mean across outstanding issues) and total bonds outstanding (in logs).

Several sample-composition considerations should be noted. First, the bond sample is restricted to firms with publicly traded debt, which tilts toward larger and more established issuers relative to the full Compustat universe. Second, the Sautner et al. (2023) climate change exposure measure, derived from earnings call transcripts, is available only from 2002 and for firms that hold such calls, introducing a selection margin that correlates with size and disclosure intensity. Third, value-weighted bond yields may change mechanically as issue composition shifts after the shock (e.g., if shorter-maturity or lower-rated bonds mature or are called). While these concerns are inherent to any firm-level bond market analysis, they caution against interpreting the bond-level heterogeneity results as precisely estimated causal effects; they are better viewed as informative cross-sectional patterns that complement the aggregate evidence.

Here I present the results from the WRDS bond market data, examining how the firm-level bond market response to CPU shocks varies with climate change exposure and financial leverage.

Baseline bond market response. Figure 6 (bond panel) shows that following a CPU shock, the average firm-level bond yield increases by approximately 20–35 basis points over quarters 2–8, peaking around quarter 7 before reverting. Log total bonds outstanding initially declines slightly on impact (approximately -0.3%), then rises by roughly 0.5% around quarters 4–6—potentially reflecting precautionary issuance or substitution from bank lending—before reverting to zero by quarter 10. Both responses are significant at the 90% level over portions of the horizon.

Climate exposure interactions. Figure 10 examines whether climate-exposed firms face differential bond market responses. The interaction of the CPU shock with climate change exposure (top row, with time FE) shows a mildly negative coefficient on bond yields (-0.02 to -0.05 percentage points) and a negative effect on log total bonds outstanding (-0.3 to -0.5%), though both are statistically imprecise at the 90% level for much of the horizon. Without time fixed effects (middle and bottom rows), the direct effect of CPU shocks on bond yields is clearly positive (20–50 basis points), confirming the aggregate spread widening documented in the monthly analysis.

Leverage interactions. Figure 11 shows that leverage interacts with CPU shocks in a manner consistent with a credit constraint mechanism. The interaction of CPU with leverage on bond yields (top-left) begins near zero but rises to approximately 5–10 basis points by quarters 6–8, indicating that more leveraged firms experience differentially larger increases in borrowing costs. More strikingly, the interaction on log total bonds outstanding (top-right) is negative and economically large, with the point estimate reaching approximately -1.5 to -2.0% over quarters 2–6. This indicates that a one-standard-deviation increase in within-firm leverage is associated with a substantial additional contraction in bond market access following CPU shocks. These results provide direct evidence that the financial channel penalizes leveraged balance sheets in the bond market.

Triple interaction: leverage, exposure, and CPU. Figure 12 reports specifications that jointly interact the CPU shock with leverage and climate change exposure. The leverage interaction on bond yields (top-left) mirrors the individual specification, with a positive coefficient of approximately 5–10 basis points at medium horizons. The climate exposure interaction on bond yields (middle-left) is close to zero. The triple interaction (bottom row) shows a modest positive effect on bond yields (roughly 5–10 basis points) and a positive effect on log total bonds outstanding (approximately 1–2%), though both are imprecisely estimated at the 90% level. The positive sign on the triple interaction for bonds outstanding may reflect the precautionary issuance motive: firms that are both highly leveraged and climate-exposed may attempt to lock in bond funding in advance of anticipated tightening.

Reconciling equity and bond market evidence. The bond market results provide an important complement to the Compustat investment analysis. While the equity-side leverage interaction shows that less-leveraged firms cut investment more (the Ottonello-Winberry channel), the bond market evidence reveals that leverage amplifies the *financial* penalty: higher-leverage firms face rising yields and declining bond volumes. The reconciliation is that the financial constraint channel operates on the cost and availability of external finance, while the investment adjustment reflects the interplay of financial flexibility and real-option considerations. Firms with lower leverage can afford to "wait and see," reducing investment as a precautionary response, while more leveraged firms—already constrained—face rising funding costs but limited room to adjust their capital spending further.

7 Mechanisms

The evidence assembled across aggregate, state-dependent, and firm-level analyses points to a multi-channel transmission mechanism. This section synthesizes the findings and relates them to the theoretical literature.

7.1 Risk Premia and Intermediary Balance Sheets

The earliest financial response to a CPU shock is a rise in the excess bond premium—the component of credit spreads that reflects the effective risk-bearing capacity of financial intermediaries (Gilchrist and Zakrajšek, 2012). The EBP increases by approximately 5–8 basis points within the first year at both monthly and quarterly frequencies, and this response is concentrated almost entirely in periods of tight financial conditions (Figure 4). Because the EBP captures fluctuations in intermediary risk appetite

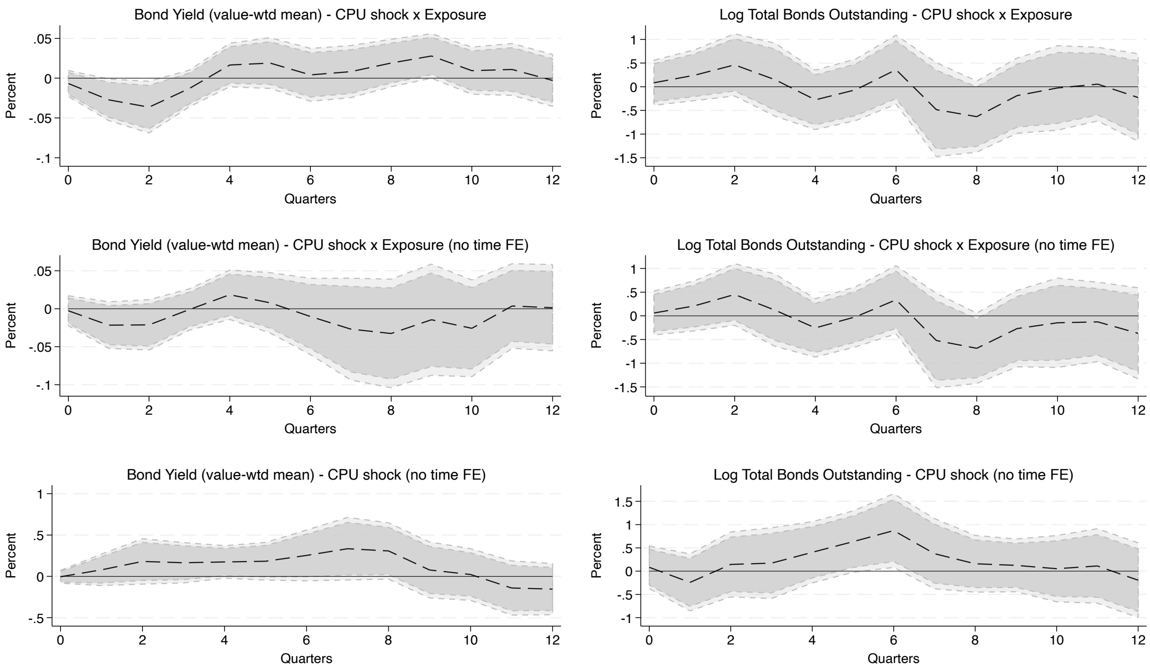

Figure 10: Bond Market Responses to CPU Shocks: Climate Exposure

Notes: Impulse responses from firm-level local projections as in equation (7). Top row: interaction with relative climate change exposure. Middle row: interaction with a brown-sector indicator. Bottom row: triple interaction. Left column: bond yield (percentage points). Right column: monthly bond return. Shaded areas denote 90% and 95% confidence bands based on standard errors clustered at the date level.

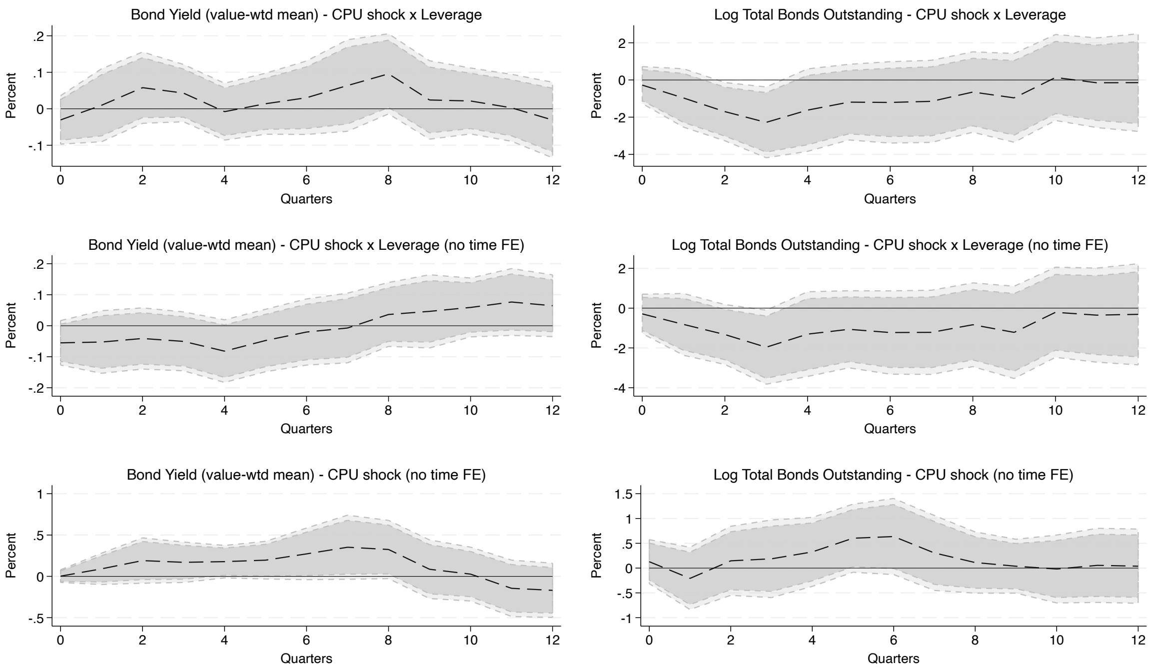

Figure 11: Bond Market Responses to CPU Shocks: Leverage

Notes: Impulse responses from firm-level local projections as in equation (7). Top row: interaction with relative climate change exposure. Middle row: interaction with a brown-sector indicator. Bottom row: triple interaction. Left column: bond yield (percentage points). Right column: monthly bond return. Shaded areas denote 90% and 95% confidence bands based on standard errors clustered at the date level.

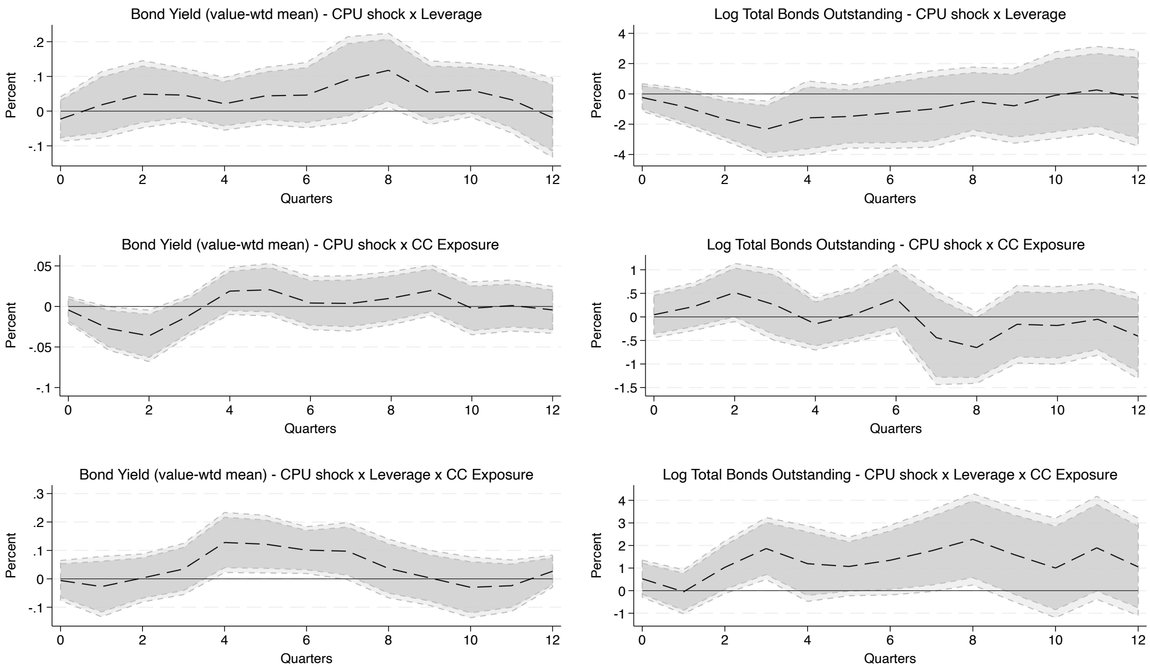

Figure 12: Bond Market Responses: Leverage and Climate Exposure Interactions
Notes: Impulse responses from firm-level local projections as in equation (7). Top row: interaction with leverage. Middle row: interaction with relative climate change exposure. Bottom row: triple interaction of leverage, climate exposure, and the CPU shock. Left column: bond yield (percentage points). Right column: monthly bond return. Shaded areas denote 68% and 90% confidence bands based on standard errors two-way clustered at the firm and date level.

rather than borrower default risk—the expected default component remains economically small and statistically insignificant—the finding points to a supply-side mechanism: CPU shocks reduce the willingness of financial intermediaries to bear risk, compressing the supply of credit.

This interpretation is consistent with the financial accelerator framework of Caldara et al. (2016), in which uncertainty shocks tighten borrowing constraints through their effect on intermediary balance sheets. It also provides direct empirical support for the asset-quality channel hypothesized in Luna-Mallea (2026): as climate policy uncertainty rises, the value of carbon-intensive (brown) capital becomes more uncertain, reducing the collateral value backing credit relationships and thereby tightening the effective lending capacity of the financial sector.

7.2 Bank Lending Standards and Credit Supply

The SLOOS evidence reveals a credit supply mechanism. Banks actively tighten lending standards for C&I loans in the first one to four quarters after a CPU shock, with the tightening concentrated among small firms—precisely the borrowers with the weakest outside options and most opaque balance sheets. The decomposition analysis in Section 5.3 shows that conditioning on observed standards absorbs the short-run decline in C&I loans, consistent with the initial credit contraction being closely associated with the tightening of bank lending standards.

The subsequent persistence of the loan contraction beyond the period of active tightening (quarters 4–12) suggests that the initial supply shock triggers a broader credit cycle. As firms face restricted credit access, they cut investment, which reduces income and deteriorates balance sheets, further reducing their creditworthiness—the classic financial accelerator loop (Gilchrist and Zakrajšek, 2012; Caldara et al., 2016). The state-dependent results reinforce this: when financial conditions are already tight, the feedback from tighter standards to reduced lending and further balance sheet deterioration is amplified.

7.3 State-Dependent Financial Amplification

The state-dependent analysis (Section 5.4) reveals that the financial transmission of CPU is highly nonlinear. During tight financial conditions:

- Investment declines are several times larger than during loose conditions, with a joint p-value for equality across states of 0.01.
- The excess bond premium rises significantly ($p(\text{tight} = 0) = 0.03$), while it remains flat under loose conditions.
- High-yield and investment-grade corporate spreads widen sharply, with the broad

U.S. corporate OAS showing equality strongly rejected.

- Loan demand drops sharply in tight states, and lending standards tighten more forcefully for small firms.

This pattern is consistent with a "doom loop" in which CPU shocks interact with pre-existing financial vulnerabilities to produce nonlinear disruption: the uncertainty shock raises risk premia, which tightens credit, which depresses investment and output, which further deteriorates balance sheets and raises risk premia. Critically, this financial-conditions state dependence is distinct from the business-cycle state dependence documented in Luna-Mallea (2026). The ANFCI strips out the cyclical component, isolating financial conditions *per se*. The finding that financial conditions—not the output gap—drive the amplification provides a more targeted test of the financial transmission mechanism.

7.4 Channel 4: Firm-Level Heterogeneity and the Role of Leverage

The firm-level evidence adds two dimensions to the mechanisms discussion. First, firms with greater climate change exposure (Sautner et al., 2023) experience differentially larger investment declines, confirming that CPU transmits with greater force to firms whose assets and business models are more vulnerable to climate transition risk. This is consistent with a revaluation channel in which markets reprice climate-exposed assets under policy uncertainty, increasing the cost of capital for these firms.

Second, the leverage interaction reveals a nuanced pattern. In the Compustat investment data, the interaction is *positive*: less-leveraged firms cut investment more than highly leveraged firms. This echoes Ottonello and Winberry (2020), who show that firms farther from their borrowing constraint exhibit stronger responses to aggregate shocks. In the bond market, however, leverage clearly amplifies the financial penalty: highly leveraged firms face rising bond yields and declining bond market access. The reconciliation is that the financial channel operates through two distinct margins—the *price* of credit (which worsens disproportionately for leveraged firms) and the *quantity* of real investment adjustment (which is larger for unconstrained firms that have the flexibility to respond). This distinction has important implications: financial stress accumulates on the balance sheets of leveraged firms even when they cannot reduce investment, potentially creating latent fragility that materializes in subsequent periods.

7.5 Reconciling Supply and Demand Interpretations

Gavriilidis et al. (2026) characterize CPU as operating like a supply shock: output falls and prices rise simultaneously. My findings are consistent with this characterization but add an important qualification. The financial transmission channel operates primarily through *demand-side* mechanisms: rising risk premia, tighter credit supply, and reduced investment, as in the theretical front in Luna-Mallea (2026). The supply-shock properties documented by GKRS (rising commodity prices and inflation) coexist with these financial demand-side effects. The relative importance of each channel depends on the state of financial conditions: when credit is tight, the financial amplification mechanism dominates, producing dynamics that more closely resemble a demand shock (synchronized declines in output, investment, and credit). When financial conditions are loose, the supply-shock characteristics dominate and financial amplification is negligible.

This state-dependent reconciliation provides a unified interpretation: CPU shocks have a supply-side origin (they raise costs and prices in carbon-intensive sectors) but their financial propagation is demand-like (they reduce credit availability and investment). The balance between these two forces varies with the state of the financial system.

8 Policy Implications

The evidence assembled in this paper carries implications for three distinct policy audiences: central banks, financial regulators, and climate policymakers.

8.1 Central Banks and Monetary Policy

The results carry two observations relevant for central bank surveillance. First, the finding that CPU shocks raise risk premia and are associated with tighter credit supply independently of the policy rate suggests that standard interest rate cuts may face limitations in offsetting the contractionary effects of climate policy uncertainty. The federal funds rate declines endogenously by 10–15 basis points following a CPU shock, yet credit spreads widen and bank lending contracts persistently despite this accommodation. This is consistent with a reduced-effectiveness dynamic (Tenreyro and Thwaites, 2016) in which monetary easing loses traction when the financial system itself is a source of amplification. In periods of tight financial conditions, when the excess bond premium spikes and lending standards tighten sharply, the estimates suggest that targeted credit facilities or macroprudential tools may be a useful complement to conventional rate cuts.

Second, the state-dependent results suggest that the monetary policy trade-off created by CPU shocks—simultaneous output declines and inflation pressure, as documented by Gavriilidis et al. (2026)—may be most acute precisely when financial conditions are already adverse. Central bank communication strategies may therefore benefit from accounting for the financial conditions backdrop when calibrating responses to climate policy announcements.

8.2 Financial Regulators and Stress Testing

Current NGFS climate stress-testing scenarios focus on physical risk and transition risk under deterministic policy paths—for instance, comparing the economic impact of an orderly versus a disorderly transition. The results in this paper suggest that these scenarios may understate financial system vulnerability by omitting the effects of policy *uncertainty*. The state-dependent evidence shows that CPU shocks are associated with excess bond premium increases of 20–30 basis points during tight financial conditions and high-yield spread widening of 80–100 basis points, magnitudes that are relevant for bank capital adequacy and portfolio risk assessment. These estimates provide a useful input for stress-testing scenarios that incorporate uncertainty or reversal in climate policy direction—a realistic characterization given the U.S. experience of repeated policy reversals across administrations.

Furthermore, the SLOOS evidence that CPU shocks are associated with disproportionate lending standard tightening for small firms highlights a distributional dimension of financial stability risk. Small firms lack access to corporate bond markets and are more dependent on bank credit, making them potentially more vulnerable to the credit supply channel documented here.

8.3 Climate Policymakers

The results are relevant for the *timing* and *credibility* of climate policy announcements. The state-dependent evidence indicates that CPU shocks coincide with the greatest credit market disruption when financial conditions are already tight. A major climate policy reversal during a period of elevated credit spreads and strained bank balance sheets could therefore be associated with a self-reinforcing cycle of rising risk premia, tighter lending, and falling investment. Conversely, the estimates suggest that implementing climate policies during periods of loose financial conditions would allow the financial system to absorb the associated uncertainty with comparatively modest disruption.

More broadly, the results speak to the economic value of policy predictability. The literature on economic policy uncertainty (Baker et al., 2016) has documented that policy uncertainty carries real costs; the contribution here is to show that *climate* policy uncertainty co-moves with additional financial costs that may amplify the real damage. Credible long-term policy frameworks—such as legislated carbon price trajectories with built-in adjustment mechanisms, or bipartisan commitments to regulatory stability—would reduce the narrative-event-driven uncertainty that generates the shocks studied in this paper. However, the reduced-form estimates here do not compare alternative policy designs or identify welfare effects, so these observations should be viewed as motivation for further research rather than as specific policy prescriptions.

9 Conclusion

This paper provides new empirical evidence on the transmission of climate policy uncertainty shocks through U.S. credit markets. Using the narrative identification strategy of Gavriilidis et al. (2026)—146 exogenous climate policy events spanning 1986–2019—I document three main findings.

First, CPU shocks are associated with significant credit market responses. At the aggregate level, a one-standard-deviation CPU shock raises the excess bond premium by approximately 5–8 basis points, widens high-yield corporate spreads by 15–25 basis points, and produces a persistent contraction in C&I bank loans of roughly 1–2% over 12–24 months. The SLOOS confirms that banks tighten lending standards in the first one to four quarters after the shock, and a decomposition exercise shows that conditioning on observed standards absorbs the short-run credit contraction.

Second, the financial effects of CPU are strongly state-dependent on prevailing financial conditions. When the ANFCI indicates tight conditions, CPU shocks are associated with investment declines several times larger than during loose conditions, significant increases in the excess bond premium, and sharp widening of corporate credit spreads. This nonlinear amplification pattern is distinct from the business-cycle state dependence documented in Luna-Mallea (2026) and provides a more targeted test of the financial mechanism.

Third, firm-level evidence reveals meaningful heterogeneity. Firms with greater climate change exposure (Sautner et al., 2023) experience differentially larger investment declines. In the corporate bond market, higher-leverage firms face rising yields and declining bond volumes, indicating that the financial channel penalizes leveraged balance sheets. However, the Compustat investment-leverage interaction reveals a nuanced Ottonello-Winberry pattern: less-leveraged firms, possessing greater financial flexibility, cut investment more sharply in response to CPU shocks. This distinction between the financial price effect (which penalizes leverage) and the real investment effect (which reflects adjustment flexibility) offers a richer view of how climate policy uncertainty propagates through the economy.

Several extensions merit future research. First, the aggregate analysis could be enriched with structural modeling that nests both the supply-shock channel of Gavriilidis et al. (2026) and the financial accelerator mechanism documented here. The DSGE framework in Luna-Mallea (2026) provides a natural starting point; the empirical estimates from this paper could discipline the calibration of the financial block. Second, the state-dependent results call for theoretical models of nonlinear financial amplification under climate policy uncertainty, extending the smooth-transition approach to a fully specified general equilibrium setting. Third, the firm-level bond market analysis could be deepened with bond-level data on issuance, maturity structure, and covenant provisions to trace the financial channel in finer detail. Fourth, the international dimension—whether CPU in the United States spills over to financial conditions in other economies—is an important open question, particularly given the global interconnectedness of bond markets and banking systems. Finally, a formal horse-race specification including the CPU shock alongside EPU, VIX, and monetary policy surprise series would further sharpen the distinctiveness claim, complementing the existing evidence from Luna-Mallea (2026) on the divergent macroeconomic dynamics under climate versus general uncertainty shocks.

The central message of the paper is that climate policy uncertainty does not merely reduce output and raise prices, as prior work has shown—it also co-moves with significant disruptions to the financial system that intermediates credit to the real economy. This financial amplification is concentrated in periods of tight financial conditions, creating the potential for self-reinforcing cycles of policy uncertainty and financial stress. Understanding and accounting for this channel is relevant for the design of climate transition policies, the calibration of central bank responses, and the development of climate-related financial stability frameworks.

References

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

Auerbach, A. J. and Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, 4(2):1–27.

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636.

Baker, S. R., Bloom, N., Davis, S. J., and Kost, K. J. (2026). Policy news and stock market volatility. *Journal of Financial Economics*, 175(1):104–187.

Basu, S. and Bundick, B. (2017). Uncertainty shocks in a model of effective demand. *Econometrica*, 85(3):937–958.

Berestycki, C., Carattini, S., Dechezleprêtre, A., and Kruse, T. (2022). Measuring and assessing the effects of climate policy uncertainty. *OECD*.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Bolton, P. and Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549.

Caggiano, G., Castelnuovo, E., and Groshenny, N. (2017). Uncertainty shocks and unemployment dynamics in U.S. recessions. *Journal of Monetary Economics*, 67:78–92.

Caldara, D., Fuentes-Albero, C., Gilchrist, S., and Zakrajšek, E. (2016). The macroeconomic impact of financial and uncertainty shocks. *European Economic Review*, 88:185–207.

Carattini, S., Heutel, G., and Melkadze, G. (2023). Climate policy, financial frictions, and transition risk. *Review of Economic Dynamics*, 51:778–794.

Diluiso, F., Annicchiarico, B., Kalkuhl, M., and Minx, J. C. (2021). Climate actions and macro-financial stability: The role of central banks. *Journal of Environmental Economics and Management*, 110:102548.

Favara, G., Gilchrist, S., Lewis, K. F., and Zakrajšek, E. (2016). Updating the recession risk and the excess bond premium. *FEDS Notes*.

Gavriilidis, K. (2021). Measuring climate policy uncertainty. Working Paper, Available at SSRN 3847388.

Gavriilidis, K., Känzig, D. R., Raghavan, R., and Stock, J. H. (2026). The macroeconomic effects of climate policy uncertainty. NBER Working Paper 34762, National Bureau of Economic Research.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.

Hsu, P.-H., Li, K., and Tsou, C.-Y. (2023). The pollution premium. *Journal of Finance*, 78(3):1343–1392.

Ilhan, E., Sautner, Z., and Vilkov, G. (2021). Carbon tail risk. *Review of Financial Studies*, 34(3):1540–1571.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182.

Jordà, Ò. (2009). Simultaneous confidence regions for impulse responses. *Review of Economics and Statistics*, 91(3):629–647.

Jordà, Ò. (2023). Local projections for applied economics. *Annual Review of Economics*, 15:607–631.

Ludvigson, S. C., Ma, S., and Ng, S. (2021). Uncertainty and business cycles: Exogenous impulse or endogenous response? *American Economic Journal: Macroeconomics*, 13(4):369–410.

Luna-Mallea, F. (2026). On the effects of climate policy uncertainty. *Economic Modelling*, 107523.

Montiel Olea, J. L., Plagborg-Møller, M., Qian, E., and Wolf, C. K. (2025). Local projections or VARs? a primer for macroeconomists. Technical report, National Bureau of Economic Research.

Noailly, J., Nowzohour, L., and Van Den Heuvel, M. (2022). Does environmental policy uncertainty hinder investments towards a low-carbon economy? *National Bureau of Economic Research, No. w30361.*

Noailly, J., Nowzohour, L., Van Den Heuvel, M., and Pla, I. (2024). Heard the news? environmental policy and clean investments. *Journal of Public Economics*, 238:105190.

Noailly, J. and Smeets, R. (2022). Directing technical change from fossil-fuel to renewable energy innovation: An application using firm-level patent data. *Journal of Environmental Economics and Management*, 112:102614.

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

Palikhe, A., Pokhrel, S., and Shrestha, P. (2024). Climate policy uncertainty and economic activity: International evidence. Working Paper.

Plagborg-Møller, M. and Wolf, C. K. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, 89(2):955–980.

Ren, X., Zhang, X., Yan, C., and Gozgor, G. (2022). Climate policy uncertainty and firm-level total factor productivity: Evidence from china. *Energy Economics*, 113:106209.

Sautner, Z., van Lent, L., Vilkov, G., and Zhang, R. (2023). Firm-level climate change exposure. *Journal of Finance*, 78(3):1449–1498.

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

Stock, J. H. and Yogo, M. (2005). Testing for weak instruments in linear IV regression. In Andrews, D. W. K. and Stock, J. H., editors, *Identification and Inference for Econometric Models*, pages 80–108. Cambridge University Press.

Tenreyro, S. and Thwaites, G. (2016). Pushing on a string: US monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, 8(4):43–74.

A Additional Results

This appendix presents robustness checks for the main results. I examine alternative state variables for the state-dependent analysis, and report additional smooth-transition specifications.

A.1 Alternative Horizon - Monthly IRF's

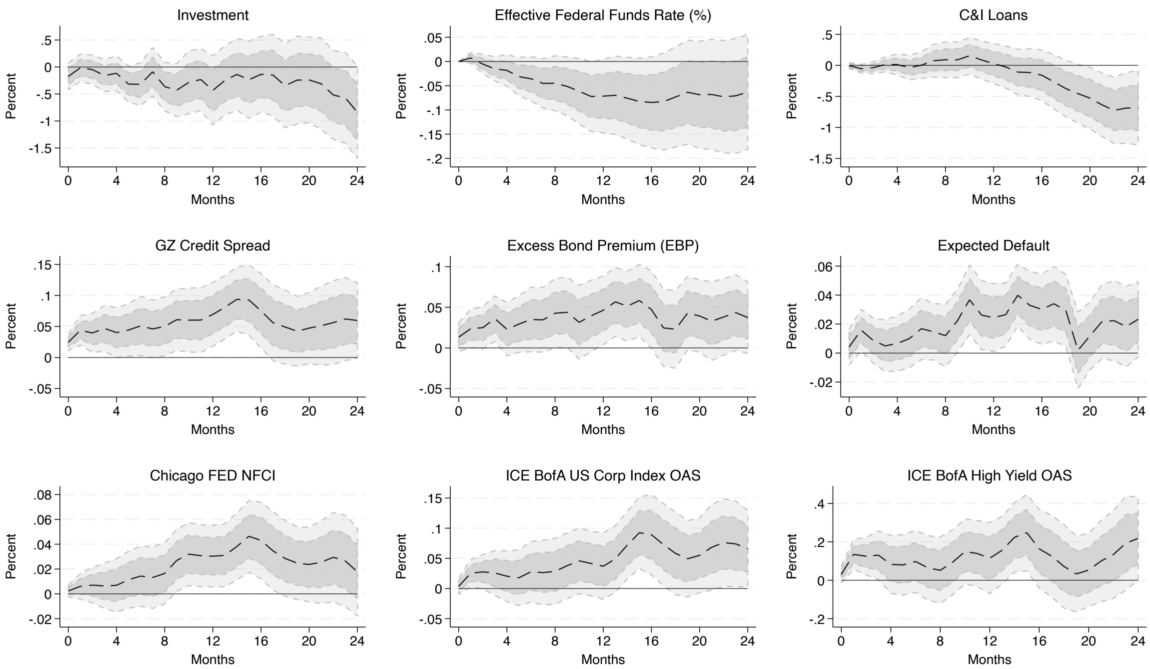

Figure 13: Aggregate Financial Responses to a CPU Shock — H=24

Notes: Impulse responses from firm-level local projections as in equation (**??**). Shaded areas denote 68% and 90%confidence bands based on robust standard errors.

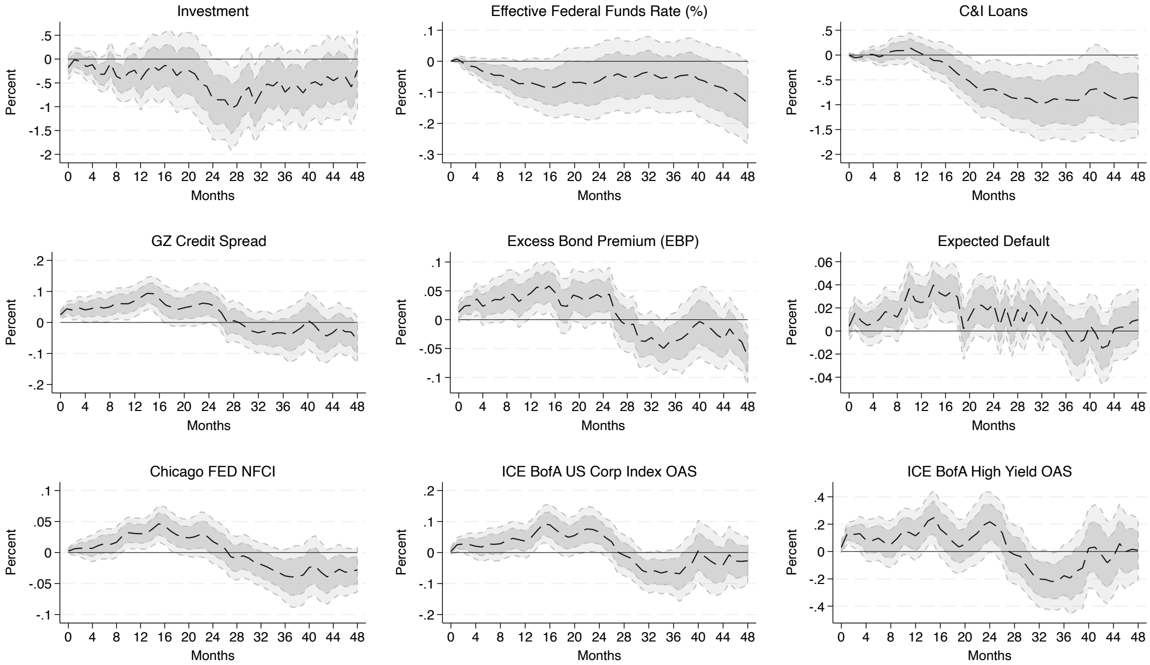

Figure 14: Aggregate Financial Responses to a CPU Shock — H=48

Notes: Impulse responses from firm-level local projections as in equation (**??**). Shaded areas denote 68% and 90% confidence bands based on robust standard errors.

A.2 Alternative Smooth Parameter - $\gamma \in \{1, 2\}$

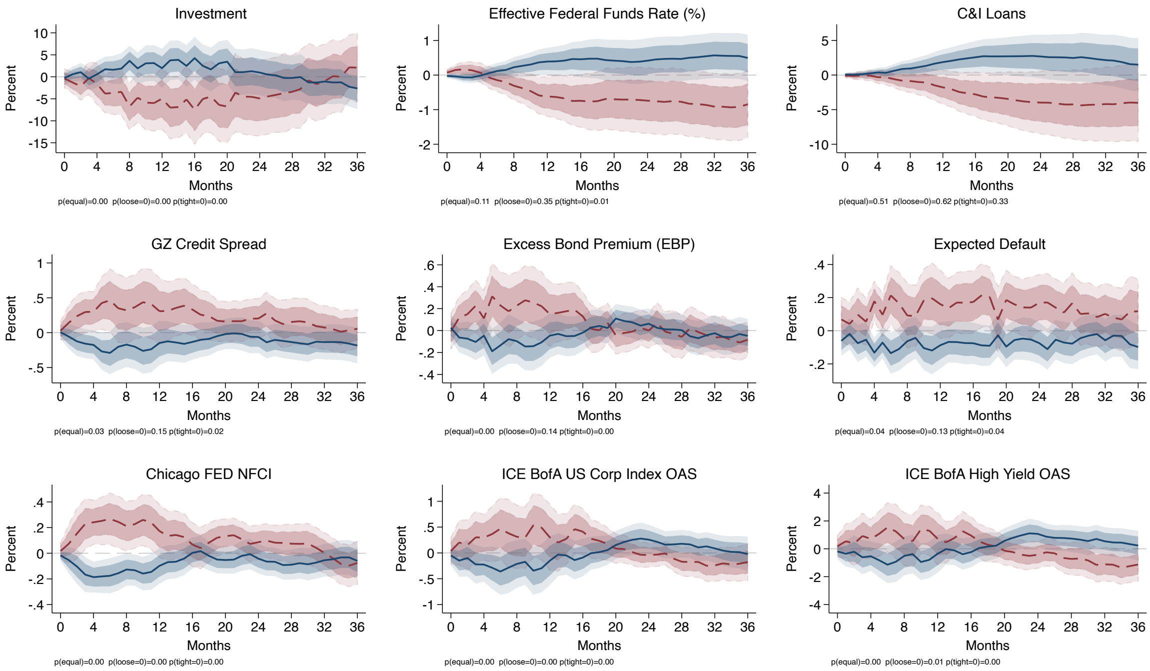

Figure 15: Monthly State Dependent IRF- State= ANFCI - $\gamma = 1$

Notes: Shaded areas denote 68% and 90% confidence bands based on standard errors clustered at the date level.

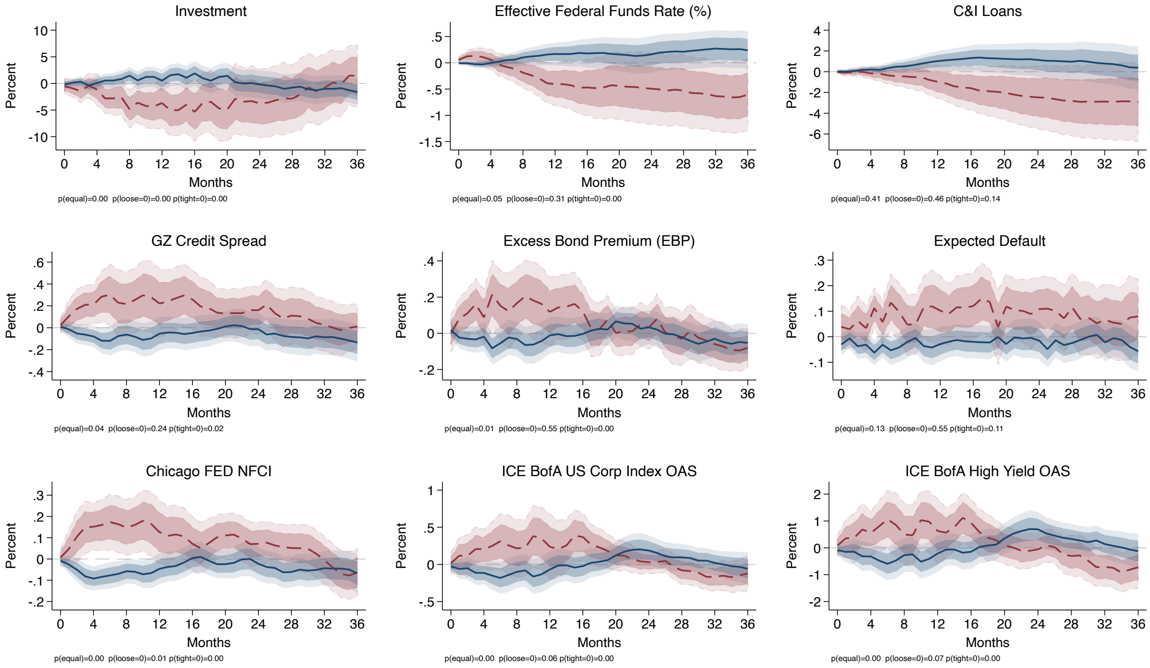

Figure 16: Monthly State Dependent IRF- State= ANFCI - $\gamma = 2$

Notes: Shaded areas denote 68% and 90% confidence bands based on standard errors clustered at the date level.

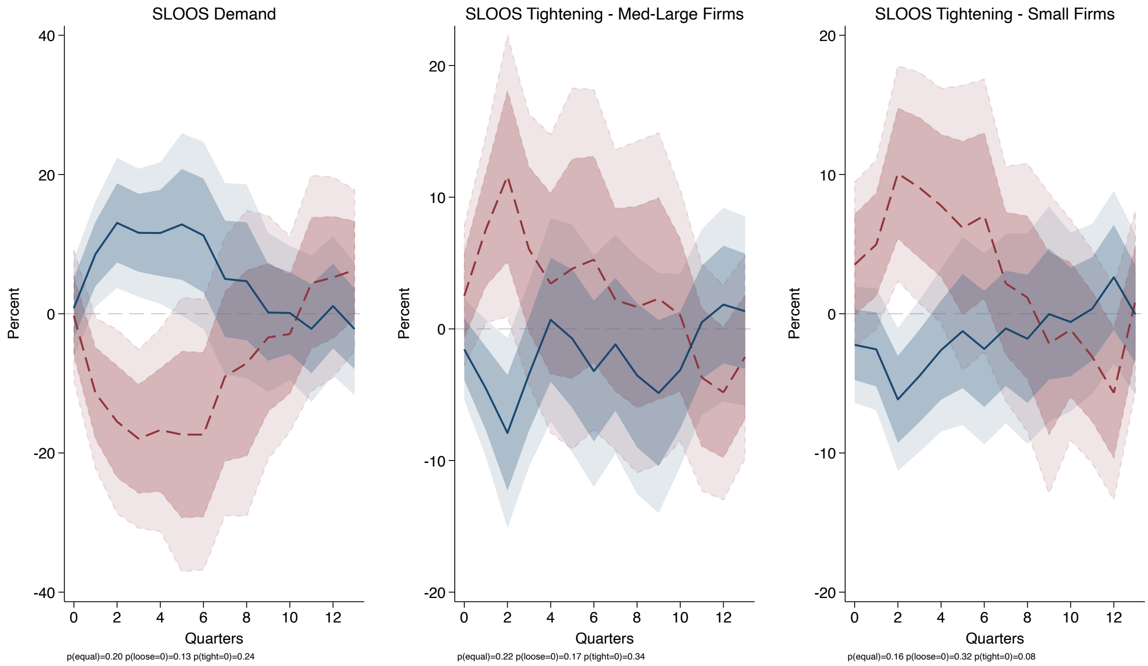

Figure 17: Quarterly State Dependent IRF- State= ANFCI - $\gamma = 1$

Notes: Shaded areas denote 68% and 90% confidence bands based on standard errors clustered at the date level.

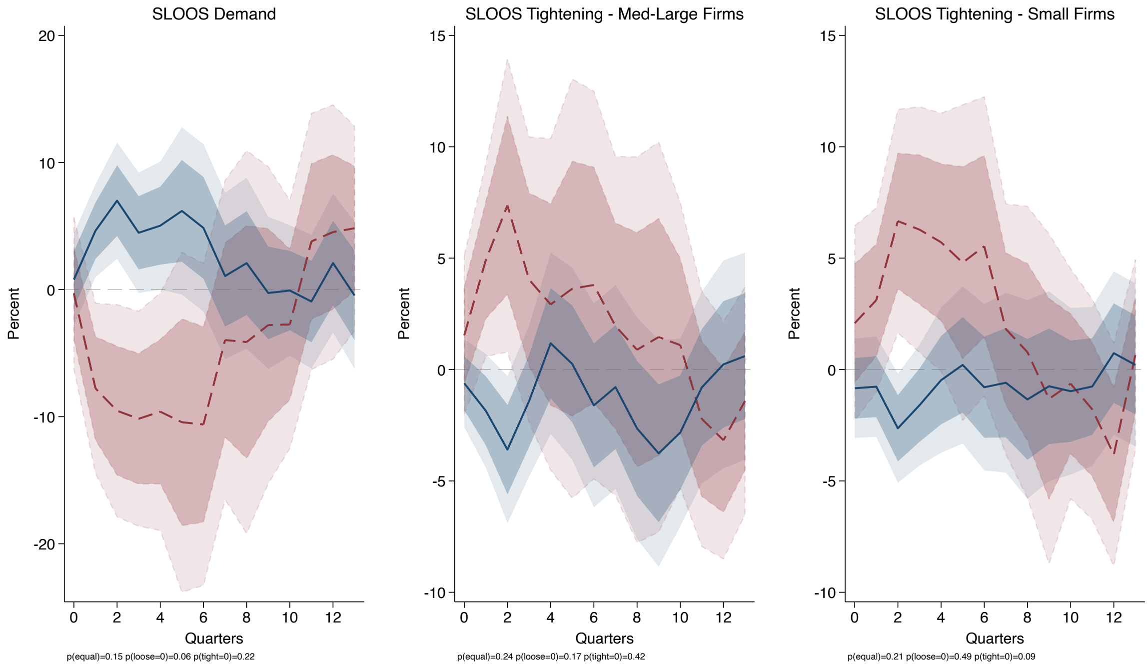

Figure 18: Quarterly State Dependent IRF- State= ANFCI - $\gamma = 2$

Notes: Shaded areas denote 68% and 90% confidence bands based on standard errors clustered at the date level.

A.3 Alternative State Variables - GZ spread

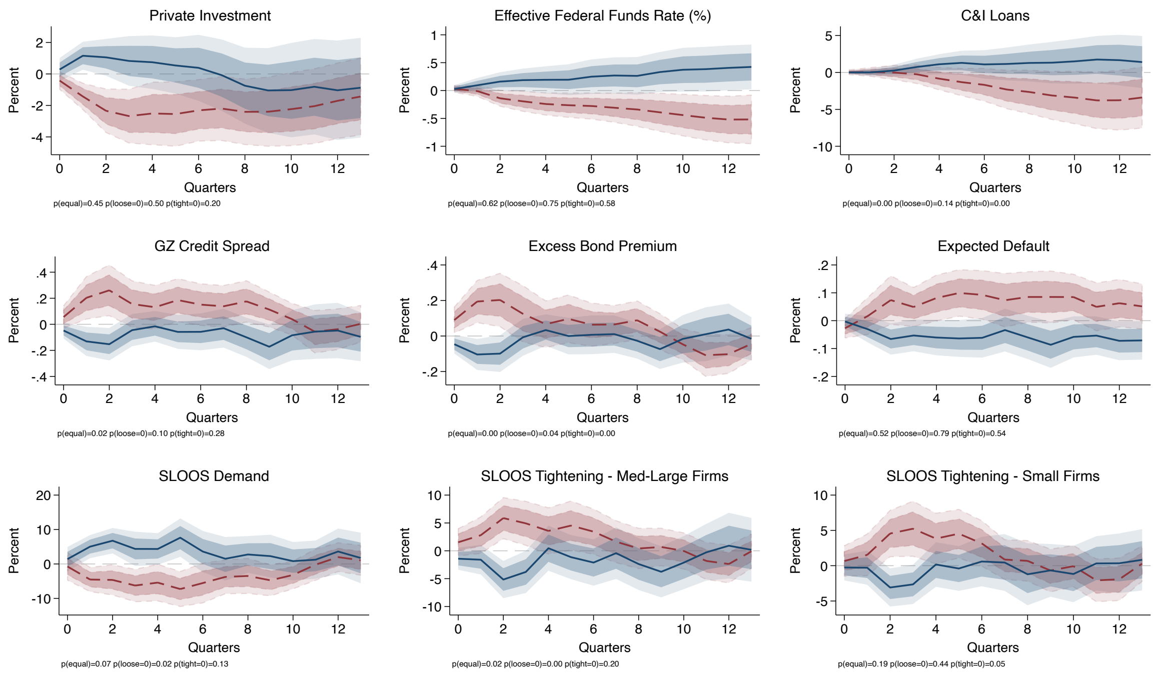

Figure 19: State Dependent IRF- State= GZ Spread

Notes: Shaded areas denote 68% and 90% confidence bands based on standard errors clustered at the date level.

B Robustness to using Luna-Mallea (2026) PCA index

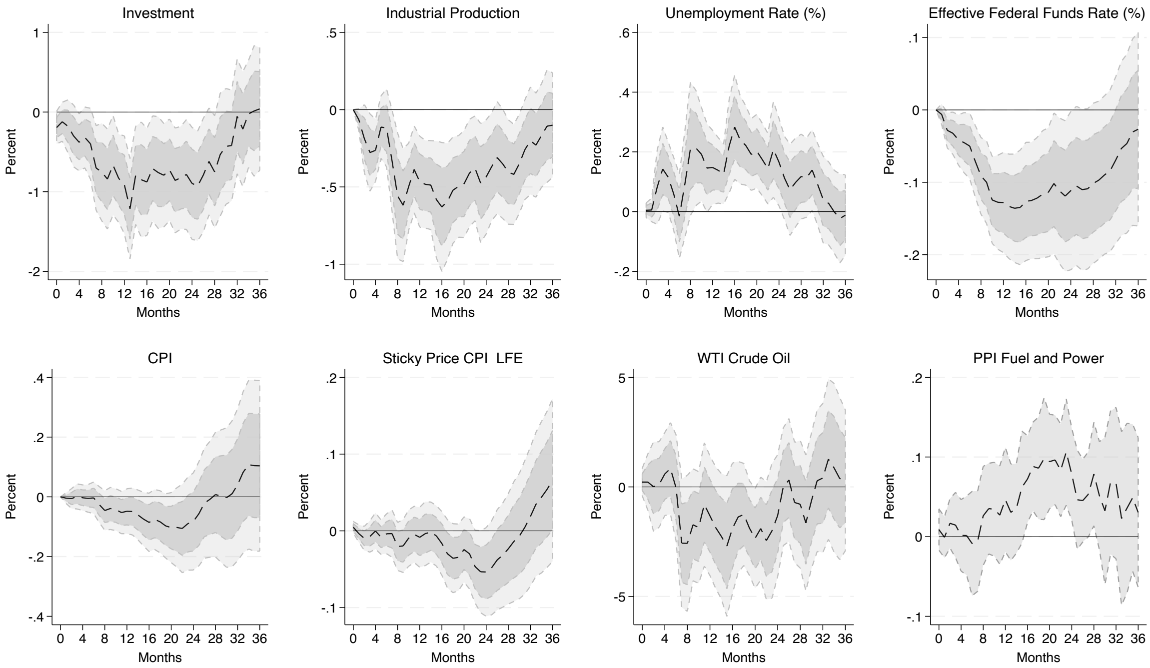

Figure 20: Aggregate Financial Responses to a CPU Shock in Luna-Mallea (2026)
Notes: Impulse responses from firm-level local projections as in equation (**??**). Shaded areas denote 68% and 90% confidence bands based on robust standard errors.

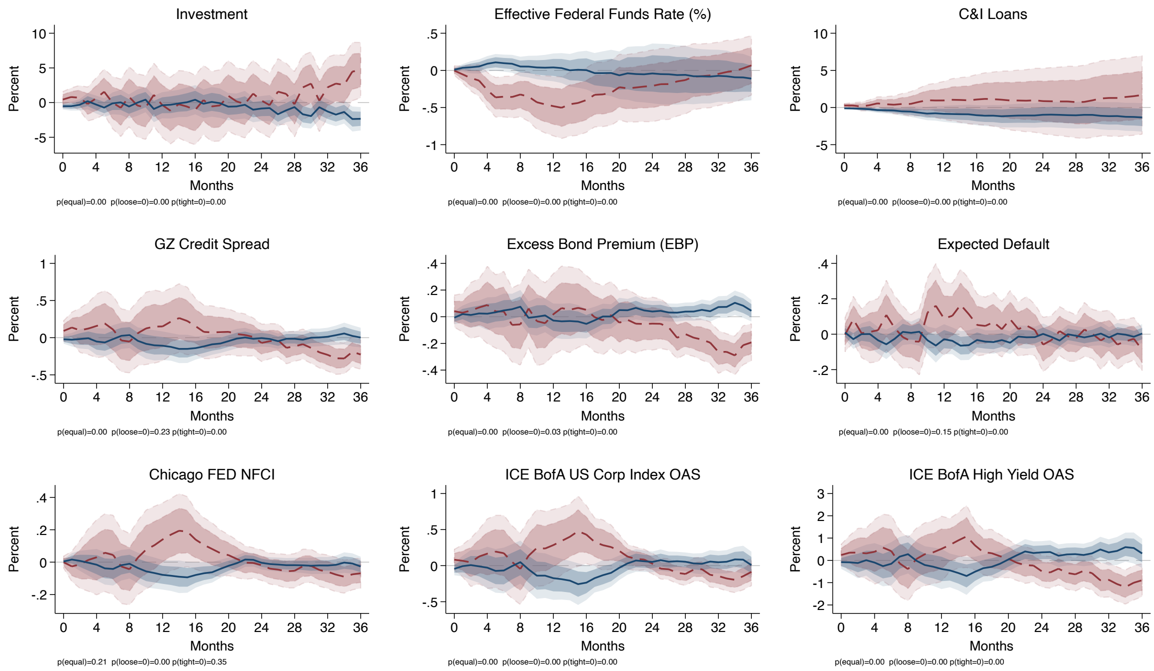

Figure 21: Monthly State Dependent IRF- State= ANFCI - $\gamma = 1.5$ - CPU Shock in Luna-Mallea (2026)

Notes: Shaded areas denote 68% and 90% confidence bands based on standard errors clustered at the date level.

C Data Sources and Construction

Table 2 summarizes the data sources for all variables used in the analysis.

Table 2: Data Sources

Variable	Source	Series / Code	Frequency
Climate Policy Uncertainty			
CPU Shock	GKRS (2026)	CPU Shock from VAR	Monthly
CPU Index (broad)	GKRS (2026)	CPU index	Monthly
Narrative Instrument	GKRS (2026)	146 events	Monthly
Financial Variables			
Excess Bond Premium	Federal Reserve Board	GZ EBP series	Monthly
GZ Credit Spread	Federal Reserve Board	GZ spread series	Monthly
C&I Loans	FRED	BUSLOANS	Monthly
SLOOS (Large firms)	FRED	DRTSCILM	Quarterly
SLOOS (Small firms)	FRED	DRTSCIS	Quarterly
SLOOS (Demand)	FRED	DRTSCLM	Quarterly
High Yield OAS	FRED	BAMLH0A0HYM2	Daily → Monthly
BBB OAS	FRED	BAMLC0A4CBBB	Daily → Monthly
IG Corp OAS	FRED	BAMLC0A0CM	Daily → Monthly
AAA OAS	FRED	BAMLC0A1CAAA	Daily → Monthly
S&P 500	FRED	SP500	Daily → Monthly
NFCI / ANFCI	FRED	NFCI, ANFCI	Weekly → Monthly
Macroeconomic Variables			
Industrial Production	FRED	INDPRO	Monthly
Unemployment Rate	FRED	UNRATE	Monthly
CPI (All Urban)	FRED	CPIAUCSL	Monthly
3-Month T-Bill	FRED	TB3MS	Monthly
Fed Funds Rate	FRED	FEDFUNDS	Monthly
VIX	FRED	VIXCLS	Daily → Monthly
EPU Index	policyuncertainty.com	U.S. EPU	Monthly
Firm-Level Data			
Compustat Fundamentals	WRDS / Compustat	fundq	Quarterly
Climate Change Exposure	OSF / Sautner et al.	CC exposure	Quarterly
Corporate Bond Data	WRDS / Mergent FISD	Bond yields, returns	Monthly

D Firm-Level Sample Construction

The firm-level panel is constructed from Compustat North America Fundamentals Quarterly, following closely the sample construction in Gavriilidis et al. (2026, Appendix C.2). The key steps are as follows:

1. **Universe:** All firm-quarter observations in Compustat North America with non-missing total assets ($\texttt{atq} > 0$) and incorporated in the United States ($\texttt{fic}$ = "USA").
2. **Industry exclusions:** Financial firms (SIC 6000–6999) and utilities (SIC 4900–4999) are excluded, as their investment behavior and leverage structures are not comparable to nonfinancial firms.
3. **Variable construction:**
   - Investment rate: $\texttt{capxy}/L.\texttt{ppentq}$.
   - Log capital expenditures: $100 \times \ln(\texttt{capxy})$.
   - R&D intensity: $\texttt{xrdq}/\texttt{atq}$.
   - Book leverage: $(\texttt{dlcq} + \texttt{dlttq})/\texttt{atq}$.
   - Cash-to-assets: $\texttt{cheq}/\texttt{atq}$.
   - Return on assets: $\texttt{oibdpq}/\texttt{atq}$.
   - Tobin's Q: $(\text{market equity} + \texttt{dlcq} + \texttt{dlttq})/\texttt{atq}$.
   - Log total assets: $\ln(\texttt{atq})$.
4. **Winsorization:** All ratios are winsorized at the 1st and 99th percentiles within each quarter to mitigate the influence of outliers.
5. **Climate exposure merge:** The Sautner et al. (2023) climate change exposure data are merged on gvkey and fiscal quarter. The merge yields approximately 84,000 firm-quarter observations with non-missing climate exposure (2002–2019).
6. **Quarterly CPU shock:** The monthly GKRS structural shock is aggregated to quarterly frequency by summing within each calendar quarter, following the convention in Gavriilidis et al. (2026).

The final firm-level sample contains approximately 196,000 firm-quarter observations spanning 1986:Q1–2019:Q4, with the climate-exposure subsample covering 2002:Q1–2019:Q4.