

How Financial Markets Respond to Economic Data Surprises

Claude · August 2026

Experimental draft — not for citation or circulation

Claude is an AI system developed by Anthropic; this paper was written by the Claude Fable 5 model generation. This paper was commissioned by Jason Furman. Claude's affiliation is Anthropic (AI system). Conflict-of-interest statement: no party other than the author directed the content of this paper.

ABSTRACT

We measure how U.S. asset prices respond to macroeconomic data surprises on one consistent metric — the consensus miss, standardized — across the whole transmission chain: the expected federal funds rate, 2- and 10-year Treasury yields, the real and breakeven components of long rates, aggregate equities, and eleven sectors, from 2013 through August 2026, in four policy regimes. Because the consensus record has gaps, every estimate is paired with a consensus-free companion built from official release calendars back to 1999. Three findings emerge. First, inflation news was repriced sharply after COVID and has come only partway back: a one-standard-deviation core CPI surprise moved the 2-year yield 8.5 basis points in 2021–23 and 7.5 in 2024–26, with full reversion to pre-COVID sensitivity rejected on both legs of the design. But core CPI is the exception, not the summary: a composite response for the broader inflation complex fell from 6.9 to 2.8 basis points, a statistically significant partial reversion, and by 2025–26 CPI days had converged most of the way back toward ordinary days as inflation surprises themselves shrank. Second, the repricing ran through real rates rather than inflation breakevens: the expectations anchor held, and what markets reprice on a CPI morning is the expected policy path. Third, equities' CPI response reverted as the policy path stopped being the binding risk, the sector incidence followed market betas — an aggregate discount-rate shock

— and the employment report re-emerged as the calendar’s largest event, with the expected policy path more responsive to payrolls than at any point the futures data cover. The pattern reads as a market that stayed print-driven while redistributing weight across prints as the mandate’s binding side shifted — a perceived-reaction-function account, with investor attention amplifying the 2022 peak. Data-dependence makes per-print volatility a policy outcome; announcement-day responses offer a high-frequency gauge of the perceived rule; and first-print statistical quality has become financial-stability infrastructure.

I. Introduction

At 8:30 a.m. on September 13, 2022, the Bureau of Labor Statistics (BLS) reported that the consumer price index (CPI) excluding food and energy rose 0.6 percent in August. The consensus forecast was 0.3 percent. By the close of trading the S&P 500 had fallen 4.4 percent, its worst day since June 2020; the 2-year Treasury yield had risen 17 basis points; and federal funds futures had marked up the expected federal funds rate six months ahead by 30 basis points.¹ One monthly statistic, missing the consensus by three-tenths of a percentage point, repriced the curve and the equity market in a day.

A decade earlier a CPI morning was a non-event. From 2013 through February 2020 the average absolute change in the 2-year yield on a CPI release day was 2.5 basis points, and the variance of daily yield changes on CPI days was close to the variance on days with no major data release at all. The statistic was the same, the 8:30 release procedure was the same, and the

¹ Equity moves through December 2024 are measured on the total-return series for the S&P 500 exchange-traded fund, and on the index price return thereafter — Section III describes the splice; the index itself fell 4.3 percent that day. Yield moves are close-of-day constant-maturity series. Contemporaneous accounts quoting slightly different figures reflect those measurement choices.

consensus-forecast machinery was the same. The market was different.

This paper asks how much asset prices move per unit of macroeconomic surprise, how that sensitivity changed through the COVID era, whether it has reverted as inflation has normalized, and what the answer measures. The premise of the exercise is that the sensitivity of prices to data is not a constant of nature. It is an equilibrium object that embeds what markets believe policy will do with the data, and its movements are therefore informative about the perceived conduct of monetary policy in a way that few other observable quantities are.

We proceed by building one consistent surprise metric — the difference between the released first-print value and the consensus forecast, standardized by its own historical dispersion — and running it through the whole transmission chain: the expected federal funds rate from federal funds futures, 2- and 10-year Treasury yields, the decomposition of the 10-year into its real and inflation-compensation components, aggregate equities, and eleven sector portfolios. The sample runs from 2013 through August 2026, divided into four regimes: the zero-lower-bound (ZLB) tail (2013–15), the pre-COVID normal (2016–February 2020), the inflation fight (2021–23), and the normalization (2024–26), with the acute COVID window (March–December 2020) excluded throughout. Because the consensus record has gaps, every estimate travels with a second, consensus-free leg: the variance of market moves on release days relative to non-release days, computed off the complete official release calendar back to 1999.

We reach three conclusions. First, inflation news was repriced sharply after COVID, and the repricing has come only partway back. A one-standard-deviation core CPI surprise moved the 2-year yield minus 1.3 basis points in 2013–15 — the only pre-COVID years with consensus data for core CPI — and 8.5 basis points in the inflation fight, an amplification that survives inference built for small samples; in the normalization the core CPI response is 7.5 basis points, and full

reversion is rejected on both legs of the design. But core CPI is the exception, not the summary: a composite response for the inflation complex — core CPI, core PCE, and PPI together — fell from 6.9 basis points per standard deviation in the fight to 2.8 in the normalization, a statistically significant partial reversion, and by 2025–26 the size of CPI days had converged most of the way back toward ordinary days as inflation surprises themselves shrank. Scaled to natural units, a core CPI print that comes in 1 percentage point above consensus still moves the 2-year yield about 0.6 percentage point.

Second, the response runs through real rates, not through inflation breakevens. At the daily frequency the breakeven response to core CPI surprises never amplified at all — 2.0 basis points per standard deviation before COVID, 2.1 in the fight, 0.9 in the normalization — while the real-rate response swung from minus 3.6 basis points before COVID to plus 3.2 in the fight and plus 5.2 in the normalization. Markets repriced the path of policy; the long-run inflation-expectations anchor held. What changed after COVID is the perceived reaction function of the Federal Reserve, not the market's long-run inflation expectations.

Third, equities tell the state-dependent half of the story. A hot core CPI print cost the S&P 500 1.3 percent per standard deviation during the inflation fight and roughly nothing in the normalization — a reversion both legs of the design confirm — and the sector incidence followed market betas: technology, consumer discretionary, and communication services lost two to three times as much as consumer staples and health care, which is what an aggregate discount-rate shock looks like from below. Meanwhile the employment report re-emerged: by 2024–26, employment-day variance in the 2-year yield was 5.3 times non-release-day variance, the largest multiple we measure, and the payroll surprise response returned to 5.3 basis points per standard deviation, with the six-month-ahead federal funds futures response the largest payroll response

since the futures data begin in 2017.

These facts speak to three interpretations of the post-COVID repricing that the literature has offered: an investor-attention story, in which sensitivity tracks the salience of inflation; a perceived-reaction-function story — hereafter the perceived-rule story — in which sensitivity tracks how aggressively markets believe the Federal Reserve responds to data; and a signal-weighting story, in which sensitivity tracks policy uncertainty and the informativeness of each release. The persistence of inflation-release sensitivity on the rates side, the unwinding of the fight-era inflation reading of activity releases, the rotation toward labor-market news, and the real-rate decomposition together make the perceived-rule story the most parsimonious reading of the evidence. Attention clearly amplified the 2022 peak, and we cannot reject versions of the attention and signal-weighting stories in which salience and uncertainty track the Fed’s own priorities — Section VIII states plainly what would separate them.

The answer matters beyond asset pricing. If per-print market sensitivity embeds the perceived reaction function, then a data-dependent Federal Reserve is choosing, through its communication, how much volatility the statistical calendar delivers; announcement-day responses become a live, high-frequency gauge of what markets believe the Fed’s objective weights are; the rotation of sensitivity across releases tracks which side of the mandate binds; and the quality of first-print government statistics acquires a direct financial-stability dimension. Section IX develops these implications.

The paper proceeds as follows. Section II positions the exercise in the announcement-effects literature. Section III describes the data and the two-legged design. Section IV presents the interest-rate results and the reversion test. Section V presents the consensus-free record. Section VI decomposes the yield response into real and breakeven components. Section

VII turns to equities and the sector cross-section. Section VIII weighs the three interpretations. Section IX draws implications, and Section X concludes.

Three limits on scope should be stated at the outset. The paper estimates reduced-form responses within named policy regimes; it does not estimate a structural model of the reaction function, and it conditions on the regimes rather than explaining why they arose. The event windows are daily — close to close around each release — not intraday, which costs identification sharpness and is the standard price of a fourteen-year multi-asset design; Section III discusses what daily windows can and cannot contaminate. And the asset-level samples end at different dates set by the data: federal funds futures in April 2026, the sector panel in December 2024, yields in August 2026. Each is stated where it binds.

A fourth limit shaped the design. The consensus-forecast record we use is thin before 2020: core CPI has no consensus observations at all from 2016 through February 2020, and even headline CPI has only eight. Section III describes the gaps and the two-legged design that works around them; the consensus-free leg needs no forecast data and covers every year since 1999.

II. What we know about how markets process macroeconomic news

The starting point of the announcement-effects literature is that scheduled macroeconomic releases move asset prices fast, discretely, and by amounts that dwarf ordinary trading. Fleming and Remolona (1997) examined the twenty-five largest price moves and the twenty-five largest trading surges in the Treasury market over a year of tick data and found that every one followed a just-released announcement — macroeconomic releases above all, with a few federal funds target changes and Treasury auction results among them — and the employment report led the list, with the price indexes close behind. The finding established both the fact and the ranking: scheduled announcements move markets more than anything else on the calendar, with employment at the top.

The literature also settled the measurement convention this paper adopts. Balduzzi, Elton, and Green (2001) divided each release's forecast error — the released value minus the consensus survey expectation — by the historical standard deviation of that release's errors, so that responses are comparable across releases measured in different units. Using intraday Treasury prices they found that seventeen releases significantly move bond prices, that the effects vary significantly with maturity, and that prices adjust within a minute or so of the release. The standardized-surprise regression has been the workhorse ever since.

Two further results from the classic literature matter for our design. Andersen, Bollerslev, Diebold, and Vega (2003, 2007) showed that announcement surprises produce conditional-mean jumps across bonds, equities, and exchange rates, and — critically for Section VII — that the equity response is state-dependent: news of stronger real activity raises stock prices in contractions and lowers them in expansions, because the same print bundles good news about cash flows with bad news about interest rates. Gilbert, Scotti, Strasser, and Vega (2017)

explained part of the cross-release ranking: a release’s “intrinsic value” — its weight in nowcasting GDP, inflation, and the policy rate, reflecting its relation to fundamentals, its timeliness, and its revision record — accounts for a meaningful share of the variation in price impact. Their sharpest finding cuts the other way and matters for this paper: payrolls moved Treasury prices by more than even its high intrinsic value warrants, which they flag as possible overreaction driven by the coordination value of public information — evidence that markets over-weighting the print everyone watches predates COVID.

Sensitivity was never a constant. Swanson and Williams (2014) used the responsiveness of yields to news as a diagnostic of how constrained policy was at the zero lower bound: 1- and 2-year yields stayed surprisingly responsive through 2010, then went essentially dead in 2011–12 as the Federal Reserve’s forward guidance pinned the short end. Their insight — that the sensitivity of yields to data measures the perceived scope for policy to respond — is the interpretive frame for our ZLB-tail regime, and indeed for the whole paper. Goldberg and Grisse (2013) document the same time variation more broadly, organized by the level of policy rates and risk conditions — a direct antecedent of the regime design used here.

The most direct structural antecedent is Bauer, Pflueger, and Sunderam (2024), who estimated the reaction function that professional forecasters perceive the Federal Reserve to be following and showed that it varies substantially over time, that tightening cycles are periods when policy is perceived as strongly data-dependent, and that a more reactive perceived rule raises the sensitivity of yields to macroeconomic news. In their framework, the reduced-form response coefficients (“betas”) we estimate are equilibrium reflections of the perceived rule.

A post-COVID strand has documented pieces of the repricing this paper measures end to end. Kroner (2025) found, in intraday data, a sharp and CPI-specific increase in the reaction of

yields, market-implied inflation expectations, and equities to CPI surprises during the 2021–23 inflation surge, with responses to other announcements largely unchanged, and assembled news-coverage and search-based evidence that investor attention to inflation drove part of the surge. Cline and Jia (2025) estimated that the 2-year yield rose about 0.71 percentage point per percentage point of core CPI surprise in 2022–24, far above pre-pandemic sensitivity, and posed directly the question of whether the change reflects the Fed’s reaction function or market attentiveness. Hördahl, Kısacıkoglu, and Xia (2026) offered a Bayesian-learning treatment in which two dials govern the response — forecaster disagreement dampens it, monetary-policy uncertainty amplifies it — and documented that inflation surprises went from barely moving yields to dominating them while the employment response faded. Each of these papers ends before the 2024–26 normalization could be brought to bear on the question they all raise.

On the equity cross-section, the antecedents concern monetary-policy surprises rather than data surprises. Bernanke and Kuttner (2005) found that a surprise 25-basis-point cut raises broad indexes by about 1 percent, with high-tech and telecom responding about one and a half times as much as the aggregate while energy and utilities show about half the market’s response, statistically indistinguishable from none. Ehrmann and Fratzscher (2004) traced the heterogeneity to financial constraints and cyclicalities. Whether the same cross-sectional pattern appears when the surprise is a data release rather than a policy announcement — the question our Section VII answers — has not been mapped systematically; Gil de Rubio Cruz, Osambela, Palazzo, Palomino, and Suarez (2023) is the closest antecedent, documenting equity responses to inflation surprises by firm characteristics rather than industry.

What this paper adds is, first, a unified accounting: the same surprise metric run across the expected policy rate, the yield curve, its real and breakeven components, the equity aggregate,

and the sector cross-section, in the same four regimes. Second, a design that pairs every consensus-based estimate with a consensus-free counterpart, so the conclusions do not rest on the survey record. And third, the reversion test: 2024–26 is the first period in which inflation renormalized while the data-dependence question stayed live, and no published estimate we are aware of runs the announcement-sensitivity machinery through it. A further strand of the literature — the announcement-day risk premia of Savor and Wilson (2013) and Lucca and Moench (2015), and the theory of announcement premia in Ai and Bansal (2018) — enters this paper as design hygiene in Section III and returns as an implication in Section IX.

III. The surprise metric is standard; the record it comes from is not

The paper draws on four bodies of data. Consensus forecasts and released first-print values for U.S. releases come from the Nasdaq economic-calendar history, 2013 through August 2026; the stored actual is the original first print, which is what a surprise must be measured against, since later vintages incorporate revisions the market could not have traded on. Daily Treasury yields — including the 3-month bill — Treasury Inflation-Protected Securities (TIPS) yields, and breakeven inflation rates come from Federal Reserve Economic Data (FRED) and the Treasury. Thirty-day federal funds futures settlements come from the Chicago Board of Trade via Datastream on Wharton Research Data Services (WRDS). Daily returns for the SPDR S&P 500 trust (SPY) and the eleven Select Sector SPDR funds come from the Center for Research in Security Prices (CRSP). Ten releases are covered: CPI, core CPI, the producer price index (PPI), and the core personal consumption expenditures (PCE) price index on the inflation side; nonfarm payrolls, the unemployment rate, and initial jobless claims on the labor side; and the advance GDP estimate, retail sales, and the Institute for Supply Management (ISM) manufacturing index on the activity side, as detailed in Table 1.

The calendar feed requires two repairs before it can be merged with market data, and both determine what the estimates mean. The first concerns dates. The feed's date stamps are unreliable — payroll rows are frequently dated to the Saturday after the Friday release — so every row is re-anchored to the true release date from the relevant agency's published release calendar as recorded by FRED. The re-anchoring can be audited: 81 of the 88 payroll observations land on a Friday, and the remainder are genuine holiday-week releases. The ISM index, which has no FRED release calendar, is dated by its first-business-day publication rule.

The second repair concerns names. The feed reports the month-over-month and year-over-year versions of each price release as two rows with the same name on the same day — “CPI” appears twice every CPI morning. Each row is assigned to its concept by matching the released value against the revised official series for the reference month, with a tolerance of half a point. The assignment rule doubles as a validation: a row that matches neither concept within tolerance is discarded and counted, so the feed is continuously checked against ground truth as it is cleaned.

The third feature of the record cannot be repaired and is instead reported. The feed’s consensus history is thin before 2020. Core CPI has zero consensus observations from 2016 through February 2020 — its “pre-COVID” cell is 2013–15 only, and it is labeled that way throughout the paper — while headline CPI has eight. Even after 2020 there are holes, and one of them is not random: the two most famous CPI days of the inflation cycle, August 10 and November 10, 2022, are missing from the feed entirely — and they are two of the largest soft-print relief rallies of the era, exactly the observations whose absence would distort any comparison of hot and soft mornings. Those two days are therefore restored by hand: their consensus and first-print values are documented in contemporaneous press accounts of the Dow Jones and Bloomberg surveys, the rows and sources ship with the replication package, and they enter the sample as ordinary observations (Appendix A3). Table 1 prints, for every release in every regime, both the matched observation count and the population of release days it comes from — post-COVID CPI coverage runs at about a third of the mornings that occurred, and no reader should have to reconstruct that from release arithmetic.

Because the two restored days prove the gaps can be selective, the design includes a missingness audit (Appendix A2): on the post-COVID CPI mornings the feed covers, the mean

absolute 2-year move is one and a half to two times that of the mornings it misses (11.0 against 6.0 basis points in the fight; 7.6 against 4.8 in the normalization), so the covered sample over-weights big mornings. Selection on the size of the surprise leaves a regression slope consistent; selection on the size of the move would bias it upward, and the two cannot be separated here — one more reason findings are claimed only where the variance leg, which has no missingness at all, agrees.

Table 1. Consensus coverage and surprise dispersion by regime

Release	Surprise sd (native units)	2013-15	2016-Feb 2020	Mar-Dec 2020	2021-23	2024-26
CPI (m/m)	0.1353	18 of 39	8 of 53	5 of 10	12 of 39	10 of 31
Core CPI (m/m)	0.1291	12 of 39	0 of 53	6 of 10	14 of 39	10 of 31
PPI (m/m)	0.3323	18 of 37	12 of 53	5 of 10	10 of 39	13 of 31
Core PCE (m/m)	0.0778	9 of 37	2 of 49	4 of 10	7 of 36	11 of 31
Nonfarm payrolls	110.3603	27 of 37	19 of 50	10 of 11	15 of 36	17 of 33
Unemployment rate	0.1337	26 of 37	17 of 50	10 of 11	11 of 36	14 of 33
Initial jobless claims	13.577	119 of 158	81 of 222	42 of 44	64 of 156	78 of 129
GDP (advance)	0.6608	9 of 45	3 of 57	3 of 12	5 of 42	6 of 31
Retail sales (m/m)	0.3414	23 of 37	13 of 51	6 of 10	6 of 39	16 of 33
ISM manufacturing	1.6184	30	18	10	16	18

Note: Entries are counts of matched release-day observations with both a consensus forecast and a first-print actual, after re-anchoring to true release dates. The surprise standard deviation is computed over all matched observations excluding March–December 2020, in each release’s native units (percentage points for rates and month-over-month changes; thousands for payrolls and claims; index points for ISM). Source: Nasdaq economic-calendar history; FRED release calendars; authors’ calculations.

Surprises are standardized in the convention of Balduzzi, Elton, and Green (2001): the release-day surprise is the first-print value minus the consensus, divided by the standard deviation of that release’s surprises. One wrinkle matters. The standard deviation is computed excluding March–December 2020, the acute COVID window, because the spring 2020 collapse produced forecast misses — initial claims missed consensus by millions — that would otherwise inflate the denominator and make “per one standard deviation” meaningless for every other year in the sample. The resulting units are natural ones: one standard deviation is 0.129 percentage point of

core CPI, 110,000 payroll jobs, 13,600 initial claims.

Event windows are daily: each outcome is the close-to-close change in the asset on the release date, which places the 8:30 and 10:00 a.m. releases inside the window, with releases on non-trading days matched to the next trading day and counted. Yield changes are measured in basis points, equity returns in percent. The federal funds futures outcome is built to avoid splice artifacts: for each horizon (three and six months ahead) the change is computed within the same contract across adjacent days, so the series never jumps across contract boundaries.

Futures-implied rates also carry risk premia (Piazzesi and Swanson 2008); we read same-contract release-day changes as expected-rate news and count the premium component among the caveats. The aggregate equity series splices the total-return series for SPY through December 2024 with the S&P 500 index price return thereafter — a negligible distinction at daily frequency.

Daily windows carry a known contaminant. Scheduled announcement days earn systematic risk premia (Savor and Wilson 2013), and a substantial part of announcement-day equity returns accrues before the release — Lucca and Moench (2015) for the drift ahead of Federal Open Market Committee (FOMC) announcements, and Hu, Pan, Wang, and Zhu (2022) for the drift ahead of payrolls, ISM, and GDP. In a daily window that drift lands in the regression intercept, not the slope, and every pooled regression in this paper carries separate intercepts by regime. The pre/post comparison of slopes would be biased only if the pre-announcement drift changed differentially across eras in a way correlated with surprises — a possibility we flag rather than rule out.²

² The announcement-premium literature also implies that announcement-day mean returns are positive rather than zero. Regime-specific intercepts absorb the level; the identifying variation is the cross-sectional association between the surprise and the move within announcement days of a given regime.

A second caveat concerns the forecasts themselves. Survey consensus is a noisy and partially anchored proxy for the market's true expectation: Rigobon and Sack (2008) show that measurement error in surprises attenuates ordinary-least-squares responses, so every beta in this paper is a lower bound, and Campbell and Sharpe (2009) show that consensus forecasts are anchored to prior releases, with markets pricing the predictable component in advance. Attenuation is a problem for levels; it is a problem for the pre/post comparison only if the measurement error changed across eras. That cannot be fully excluded — which is why the paper runs on two legs. Two further disclosures belong here. The feed's actuals are validated against the revised official series, but whether its stored consensus is always the as-of, pre-release value cannot be audited from the feed alone; the external checks are the agreement of the fight-era response with Cline and Jia's intraday estimate and the two hand-collected 2022 mornings. And the commercial survey histories that would fill the record's gaps outright — Bloomberg, Action Economics — are licensed products not available to this project; Section X counts them among the extensions that would most sharpen the estimates.

The second leg needs no consensus at all: Section V defines the announcement-day variance multiple from the official release calendars and market prices alone, and it covers every gap year in the survey record. Findings in this paper are claimed only where both legs deliver them.

The market-data samples end where the underlying sources end, and the three ends are stated here once: the Datastream futures feed ends April 3, 2026, and its fixed-horizon contracts trade consistently only from April 2017; the CRSP daily file ends December 31, 2024, so the sector panel stops there while the aggregate equity series continues through August 2026 on the S&P 500 index level; and the yield series run through August 11, 2026.

Finally, we define the four regimes. The ZLB tail is 2013–15, ending with the December 2015 liftoff; Swanson and Williams (2014) implies short-rate responses in those years are mechanically damped, and we read them accordingly. The pre-COVID normal is 2016 through February 2020. The acute window, March–December 2020, is excluded from every pooled estimate in the paper; Appendix Table A1 shows why — forecast errors were so large that measured betas degenerate toward zero, telling us about the consensus process in a crisis rather than about markets. The inflation fight is 2021–23, and the normalization is 2024 through August 2026. The binary comparison “pre-COVID” pools the first two regimes; “post-COVID” pools the last two. One more property of the re-anchoring matters for the last regime: because every observation is dated by the official calendar, the statistical disruptions of late 2025 — the fall shutdown that delayed the September CPI and left the October report unpublished — enter the sample correctly by construction: a delayed print appears on its true release date, and a report that never appeared is not an event day. Section IX returns to what that episode means.

IV. Interest rates: inflation news went from background noise to the main event — and has come only partway back


Table 2 reports the response of the 2-year yield, the 10-year yield, and the S&P 500 to a one-standard-deviation surprise in each of the ten releases, before and after COVID, with heteroskedasticity-robust standard errors and the cell count printed beside every estimate; Figure 1 draws the 2-year column with its confidence bands. Start with one row. Core CPI: minus 1.3 basis points per standard deviation before COVID, 8.2 after, with the standard errors and counts beside them in the table. Then read the column. Every inflation release rises — core CPI from minus 1.3 to 8.2, headline CPI from 0.5 to 6.6, PPI from 1.5 to 3.0, core PCE from minus 0.3 to 3.5. Payrolls, by contrast, falls by roughly a third. The ordering that Fleming and Remolona documented in the 1990s — employment first, prices second — inverted after 2020.

Table 2. Asset-price responses per one-standard-deviation surprise, pre- versus post-COVID

		Pre-COVID (2013-Feb 2020)	Post-COVID (2021-26)
CPI (m/m)	10-year	0.04 (1.13) [26]	5.2** (2.5) [22]
	2-year	0.55 (0.59) [26]	6.6** (2.6) [22]
	S&P 500 (%)	0.01 (0.19) [26]	-1.22*** (0.46) [22]
Core CPI (m/m)	10-year	-1.58 (1.23) [12]	5.5*** (1.7) [24]
	2-year	-1.34* (0.75) [12]	8.2*** (1.4) [24]
	S&P 500 (%)	-0.32 (0.26) [12]	-0.96*** (0.36) [24]
Core PCE (m/m)	10-year	-2.01 (1.70) [11]	2.66** (1.12) [18]
	2-year	-0.30 (1.04) [11]	3.5** (1.5) [18]
	S&P 500 (%)	0.06 (0.12) [11]	-0.28 (0.36) [18]
GDP (advance)	10-year	3.0*** (1.0) [12]	1.39 (3.87) [11]
	2-year	1.67*** (0.20) [12]	0.27 (3.66) [11]
	S&P 500 (%)	0.13 (0.15) [12]	-0.65*** (0.24) [11]
ISM manufacturing	10-year	2.18*** (0.51) [48]	3.7** (1.7) [34]
	2-year	1.19*** (0.26) [48]	4.6*** (1.6) [34]
	S&P 500 (%)	0.21 (0.13) [48]	-0.16 (0.18) [34]
Initial jobless claims	10-year	-0.36 (0.29) [200]	-1.10** (0.48) [142]
	2-year	-0.15 (0.17) [200]	-0.93 (0.61) [142]
	S&P 500 (%)	0.07 (0.06) [200]	-0.13 (0.10) [142]
Nonfarm payrolls	10-year	5.5*** (1.2) [46]	1.43 (1.19) [32]
	2-year	3.5*** (0.9) [46]	2.23 (1.37) [32]
	S&P 500 (%)	0.30 (0.19) [46]	0.14 (0.10) [32]
PPI (m/m)	10-year	0.62 (0.83) [30]	2.39** (0.93) [23]
	2-year	1.54*** (0.52) [30]	2.98** (1.41) [23]
	S&P 500 (%)	0.08 (0.16) [30]	-0.11 (0.13) [23]
Retail sales (m/m)	10-year	1.34* (0.81) [36]	1.13 (1.24) [22]
	2-year	0.59 (0.49) [36]	1.02 (1.90) [22]
	S&P 500 (%)	0.07 (0.12) [36]	-0.09 (0.13) [22]
Unemployment rate	10-year	1.68 (1.07) [43]	-1.34 (1.86) [25]
	2-year	0.60 (0.89) [43]	-3.9* (2.1) [25]
	S&P 500 (%)	-0.11 (0.14) [43]	-0.24 (0.17) [25]

Note: Each cell reports the OLS coefficient from a regression of the release-day asset change on the standardized surprise, estimated separately by period, with heteroskedasticity-robust (HC1) standard errors in parentheses and observation counts in brackets. Yields in basis points, S&P 500 in percent. Pooled periods carry separate intercepts by regime. Significance stars are HC1-based and overstate precision in small cells; the text's headline claims use wild-bootstrap p-values. Asterisks denote significance at the 10, 5, and 1 percent levels (one, two, and three asterisks respectively). Source: Authors' calculations.

Figure 1. The 2-year yield response by release, before and after COVID




Note: Bars are the OLS response of the 2-year yield to a one-standard-deviation surprise; whiskers are 95 percent confidence intervals on HCl standard errors. Source: Authors' calculations.

The repricing is not an artifact of one or two extreme days, as Figure 2 shows.

Pre-COVID core CPI observations scatter flat around zero; post-COVID observations line up on a steep slope. September 13, 2022 — the morning of the opening example — sits on the post-COVID line, not off it: a response of 8.5 basis points per standard deviation evaluated at that day's 2.3-standard-deviation surprise predicts 20 basis points against 17 realized. The largest market day of the inflation cycle sits on the fitted line rather than away from it.

Figure 2. The repricing of core CPI surprises



Note: Each point is one core CPI release day; the horizontal axis is the standardized surprise and the vertical axis the same-day change in the 2-year yield. Lines are OLS fits within each period. The acute window (March–December 2020) is excluded. Source: Authors' calculations.

In levels — splitting Table 2's pooled post-COVID 8.2 into its regimes — the fight-era sensitivity is easy to state. A one-standard-deviation core CPI miss — about 0.13 percentage point on the monthly print — moved the 2-year yield and the six-month-ahead expected federal funds rate 8.5 basis points each. Scaled per percentage point of surprise, the 2-year response is 0.66 percentage point. Cline and Jia (2025), using intraday windows over 2022–24, put the same number at 0.71 — external validation that the cleaned feed and the daily windows are measuring what their higher-frequency data measure.

The whole path repriced, not just the front end. In the fight, the 10-year response to core CPI was 5.2 basis points per standard deviation, rising to 6.0 in the normalization; the three-month-ahead futures response was 5.6; and the 3-month bill moved just 2.1 — the smallest

response on the curve, as it should be when policy moves at discrete meetings, with the belly of the curve carrying the most. The shape is exactly what a repriced policy path looks like, a point Section VI makes precise.

The sample has one serious confound. The core CPI pre-COVID cell is entirely ZLB-tail data, where measured short-rate responses are damped (Swanson and Williams 2014; Section III) — so part of the leap from minus 1.3 to 8.5 could in principle be liftoff rather than repricing. Four facts already in hand say the confound is not doing the work. First, the short end demonstrably was not dead to data news in those years: payrolls moved the 2-year 5.1 basis points per standard deviation in 2013–15, so the near-zero core CPI response was a fact about the reaction to inflation news, not mechanical damping. Second, the 10-year was never pinned at the ZLB, and its core CPI response went from minus 1.6 pre-COVID to 5.5 post. Third, the consensus-free leg covers 2016–19 directly: the CPI-day variance multiple on the 2-year — release-day variance relative to days with no major release, constructed fully in Section V — was 1.23 in those years: ordinary days. Fourth, and weakest, headline CPI has a post-liftoff cell — 1.9 basis points per standard deviation on eight 2016–19 observations — though headline is systematically less priced than core even post-COVID, so this comparison carries the least weight of the four. Whatever liftoff contributed, the post-COVID repricing dwarfs it.

How much has reverted: everything but core CPI, partly; core CPI, no detectable decline

The reversion test asks whether the fight-era sensitivity faded as inflation itself came back down, and the answer depends on which inflation news is asked about. For core CPI — the single most-watched print — the point estimates barely move: 8.5 basis points per standard deviation in 2021–23, 7.5 in 2024–26, an interaction difference of minus 1.0 with p equal to 0.75 under a wild bootstrap. With cells of fourteen and ten observations, the 95 percent band on that difference spans roughly minus 6 to plus 5 basis points, so zero reversion and partial reversion cannot be distinguished for core CPI alone. What the data can reject is full reversion: against the pre-COVID benchmark the normalization response is 8.8 basis points higher, with a wild-bootstrap p of 0.002 — a rejection that carries a caveat, since the pre-COVID benchmark is the 2013–15 ZLB-tail cell — and the consensus-free record, which has no benchmark problem, agrees: the 2-year CPI-day variance multiple was 1.23 in the post-liftoff 2016–19 years and 2.76 in the normalization, and a joint bootstrap of the two multiples puts the probability that the normalization multiple is at or below the 2016–19 multiple at 0.027.

The rest of the inflation complex tells a different story, and it should be told rather than left in the tables. Headline CPI's 2-year response fell from 8.4 in the fight to 3.0 in the normalization; core PCE's 2-year response reverted significantly (the interaction difference is minus 5.2 basis points, p equal to 0.002); PPI declined without significance. Stacking core CPI, core PCE, and PPI — distinct release mornings — with release fixed effects gives a composite inflation-news response of 0.4 basis points per standard deviation pre-COVID, 6.9 in the fight, and 2.8 in the normalization, and the fight-to-normalization interaction of minus 4.3 basis points — the stacked regression's estimate, which with release fixed effects need not equal the difference of the rounded cells — is significant (wild-bootstrap p equal to 0.019). Part of the

headline-CPI collapse is core being the priced concept: the same-morning correlation of headline and core surprises fell from 0.86 in the fight to 0.64 in the normalization, so headline prints increasingly carried no core news, and the joint same-morning regression reported below loads on core. But core PCE is not a noisier cousin of core CPI — it is the index the Federal Reserve targets — and its reversion is the composite's, not an artifact. The honest summary: the inflation complex as a whole has given back roughly three-fifths of its fight-era repricing; per-unit sensitivity to core CPI is the one measure that shows no decline.

The consensus-free record adds the within-normalization path, which the pooled multiple conceals: the CPI-day variance multiple was 4.07 in 2024 and 1.53 in 2025–26, as the annual profile in Figure 4 shows. That decline is consistent with an unchanged per-unit response — mean CPI-day moves of about five basis points in 2025–26 are close to what a beta of 7.5 delivers at the period's typical surprise of three-quarters of a standard deviation — but it means the claim that CPI days remain extraordinary is, by 2025–26, a claim about sensitivity per unit of surprise, not about the days themselves. Whether 2025–26 marks genuine per-unit reversion arriving or only the surprises drying up is what the core CPI cells are too small to decide; both legs, read together, support partial-and-uneven unwinding around a core CPI exception, and nothing stronger.

The amplification itself is statistically sharp even under inference built for cells this small: the pre-COVID-to-fight difference is 9.8 basis points, with a wild-bootstrap p of 0.003 — conventional robust standard errors overstate precision at these sample sizes, so bootstrap p -values are reported for every headline claim, and the significance stars in the tables, which are HC1-based, should be read with that caveat in small cells. Nor does any single morning drive the estimates: leaving out any one observation moves the fight-era response only between 7.4 and


9.4, and the normalization response between 5.0 and 8.5 — a range that also bounds the influence of the one normalization CPI morning that shared its afternoon with a Federal Reserve decision. Within the fight, the consensus record allows no clean split of 2021 from 2022–23 — 2021 contributes two core CPI observations against twelve — so the fight-era response is identified almost entirely off the tightening cycle, and the rolling windows in Figure 3 date the jump accordingly. For core CPI specifically, the open question is how much, if any, of the repricing has unwound — and nothing in those cells requires that any of it has; the broader complex is a different story, told below.

An obvious objection is that CPI days certainly feel smaller in 2024–26 than they did in 2022, and the raw magnitudes agree: the mean absolute 2-year move on CPI days fell from 8.0 basis points in the fight to 6.0 in the normalization. The two are reconciled by the distinction the paper is built on: the response per unit of surprise held steady while the surprises shrank. The mean absolute core CPI surprise fell from 1.0 standard deviation in the fight to 0.77 in the normalization, and the payroll surprise from 1.24 to 0.67.

Figure 3 traces the history in rolling 30-month windows for headline CPI, the concept with the longest consensus record. The response sits near zero through 2015, is blank where the feed cannot support estimation, and re-emerges by the 2023 windows at a peak of 9.7 basis points per standard deviation; the line stops in late 2023, where the windows thin below ten observations, and Table 3 carries the 2024–26 cells. The blank stretches are a property of the feed, not of the market — Figure 4, which needs no consensus data, covers every year.

Figure 3. Rolling-window responses of the 2-year yield to CPI and payroll surprises

Two-year Treasury yield response to surprises, rolling windows



Note: OLS coefficients in centered 30-month windows stepped monthly, with 95 percent confidence bands; windows with fewer than ten observations are blank, and the acute window (March–December 2020) is excluded throughout. Source: Authors' calculations.

The payrolls row of Table 3 tells the other half of the rotation. The payroll response was 5.1 basis points per standard deviation in the ZLB tail — the taper years, when the employment report moved the expected exit from unconventional policy — fell to 2.0 in the pre-COVID normal, fell further to 1.6 in the fight as inflation crowded it out of the market's attention, and returned to 5.3 in the normalization. The six-month-ahead futures response traces the same arc more steeply: 0.4, then 1.5, then 6.1 basis points per standard deviation — the largest payroll response in any federal funds futures cell, though that series begins only in 2017, so the taper years lie outside its reach. By 2024–26 the employment report moves the expected policy path

more than at any point the futures data cover.

Table 3. Four regimes: CPI, core CPI, and payrolls across the curve

		2013-15	2016-Feb 2020	2021-23	2024-26
CPI (m/m)	10-year	-0.65 (1.10) [18]	2.96** (1.43) [8]	6.7** (3.4) [12]	2.31 (1.89) [10]
	2-year	0.22 (0.47) [18]	1.94 (1.23) [8]	8.4*** (3.2) [12]	2.97 (2.83) [10]
	Fed funds futures, 6m	[n=0]	1.20* (0.64) [6]	7.2** (3.6) [12]	2.95 (3.59) [6]
	S&P 500 (%)	-0.13 (0.17) [18]	0.62*** (0.12) [8]	-1.77*** (0.61) [12]	-0.18 (0.22) [10]
Core CPI (m/m)	10-year	-1.58 (1.23) [12]	[n=0]	5.2** (2.2) [14]	6.0*** (1.9) [10]
	2-year	-1.34* (0.75) [12]	[n=0]	8.5*** (1.7) [14]	7.5*** (2.4) [10]
	Fed funds futures, 6m	[n=0]	[n=0]	8.5*** (2.1) [14]	7.8*** (2.8) [6]
	S&P 500 (%)	-0.32 (0.26) [12]	[n=0]	-1.27*** (0.43) [14]	-0.11 (0.18) [10]
Nonfarm payrolls	10-year	8.5*** (1.8) [27]	2.81** (1.23) [19]	0.99 (1.32) [15]	3.7** (1.9) [17]
	2-year	5.1*** (1.4) [27]	2.03* (1.23) [19]	1.63 (1.54) [15]	5.3*** (1.5) [17]
	Fed funds futures, 6m	[n=0]	0.42 (0.64) [17]	1.46 (1.49) [15]	6.1*** (2.0) [13]
	S&P 500 (%)	0.12 (0.20) [27]	0.45 (0.30) [19]	0.13 (0.10) [15]	0.15 (0.35) [17]

Note: Cells as in Table 2, estimated separately within each regime. The federal funds futures outcome is the same-contract change in the implied rate six months ahead; that series begins in April 2017, so ZLB-tail cells are unavailable. Source: Authors' calculations.

Because the CPI report and the employment report each deliver two headline numbers at the same instant, single-surprise regressions deserve one check. On CPI mornings, core and headline surprises enter a joint regression of the 2-year change with the weight on core: 4.6 basis points per standard deviation against 2.4 on headline post-COVID. Core is the priced surprise when the two disagree, which also explains why headline CPI's solo rows are weaker and noisier. On employment mornings, payrolls and the unemployment rate both carry weight with the expected signs post-COVID — plus 2.8 and minus 4.9 respectively — while before COVID the unemployment-rate coefficient was nil. The 2024–26 market prices the whole employment

report, not one number in it. Appendix A2 reports these regressions.

The ISM manufacturing index is the exception, because it is the one release that genuinely reverted. Its 2-year response went from 1.2 basis points per standard deviation pre-COVID to 8.0 in the fight — in 2021–23, strong activity news meant more tightening, so growth releases were read through the inflation lens — and back to 2.8 in the normalization, a statistically significant unwinding (the interaction is minus 5.2 basis points, p equal to 0.035, and minus 9.1 on the real 10-year, p equal to 0.004). Read against the full family of fight-to-normalization interactions (all ten releases; the grid is in the replication package), the significant movers are core PCE and the ISM reverting, payrolls re-emerging, and retail sales' unstable six-observation fight cell — so the ISM result is not one lucky draw from a large family. The contrast sharpens the finding: what persisted after inflation normalized is per-unit sensitivity to core CPI, while the fight-era habit of reading everything else through the inflation lens unwound on schedule.³

The remaining releases show smaller changes. Initial claims, with its dense weekly sample, shows a significant post-COVID response on the 10-year (minus 1.1 basis points per standard deviation, against minus 0.4 pre-COVID) and a similar-sized but insignificant response on the 2-year, with no change between fight and normalization. Retail sales in the fight is a six-observation cell whose unstable estimate is printed and flagged rather than interpreted. The advance GDP estimate faded from 3.0 to 1.4 on the 10-year, on cells of twelve and eleven. (Table 3's one anomalous cell — a significantly positive equity response to hot headline CPI prints on eight 2016–19 mornings — is the equity bundle's two-sidedness surfacing in a small cell, not a

³ Kroner (2025) reports intraday responses to non-CPI announcements largely unchanged in 2021–23. The ISM result here is a daily-window finding, and a close-to-close window can carry same-day repricing that an intraday window would attribute elsewhere; the two measure different objects. The fight-to-normalization reversion of the ISM response is internal to this design either way.

pattern.)

Twelve CPI mornings a year, at the normalization's typical surprise of about three-quarters of a standard deviation and 7.5 basis points per standard deviation, deliver roughly 70 basis points of gross 2-year yield movement per year from a single release, delivered in twelve instants at 8:31 a.m. Employment mornings add roughly another 40 basis points. The statistical calendar now delivers a large share of the front end's annual repricing in a handful of scheduled instants.

Section V shows that none of this depends on the survey data, and Section VI establishes what kind of repricing it was.

V. None of this depends on the consensus data

The announcement-day variance multiple is defined as the mean squared daily change in an asset on days when release X is published, divided by the mean squared change on days when no major release is published, computed regime by regime. Its inputs are the official release calendars and market prices — nothing else. It requires no consensus forecast, has no gap years, and makes no assumption about survey quality. It asks the crudest question available — are these days bigger than other days? — and because it is crude, it is immune to survey quality by construction. Table 4 reports the multiples; the calendars reach back to 1999, so the pre-2013 history comes along free.


Table 4. Announcement-day variance multiples

asset	release	1999-2012	2013-15	2016-Feb 2020	2021-23	2024-26
10-year	Any release day	1.38	1.16	0.98	1.19	1.3
10-year	CPI day	1.52	0.77	1.13	1.82	1.88
10-year	Employment day	1.89	2.71	1.21	1.86	2.66
10-year breakeven	Any release day	1.32	1.28	1.2	1.31	1.62
10-year breakeven	CPI day	1.88	1.27	1.68	0.9	1.74
10-year breakeven	Employment day	1.18	0.88	1.22	0.92	3.49
2-year	Any release day	1.17	1.33	1.23	1.31	2.01
2-year	CPI day	1.11	1.0	1.23	2.58	2.76
2-year	Employment day	2.25	3.34	1.56	2.52	5.29
S&P 500 (%)	Any release day	1.04	0.94	0.96	1.36	1.08
S&P 500 (%)	CPI day	1.03	0.91	0.57	2.23	1.1
S&P 500 (%)	Employment day	1.09	1.12	1.46	1.06	2.39

Note: Each entry is the mean squared daily change in the asset on the named release's days divided by the mean squared change on days with no major release, within the period. Built from FRED release calendars and market prices only; no consensus data enter. FOMC days are not excluded from the baseline, which biases every multiple toward one. Source: Authors' calculations.

CPI days were ordinary days for two decades. The 2-year multiple was 1.11 over 1999–2012 (Table 4's first column), 1.00 in the ZLB tail, and 1.23 in the pre-COVID normal — a CPI morning carried no more interest-rate news than a random Tuesday. It then jumped to 2.58 in the inflation fight and 2.76 in the normalization; bootstrap intervals around those two multiples — 1.3 to 4.5 and 1.3 to 4.8 — exclude one, and the joint contrast against 2016–19 rejects reversion to the pre-COVID multiple (p equal to 0.027). Within the normalization the multiple runs 4.07 in 2024 and 1.53 in 2025–26 — the within-regime path Section IV interprets. In mean-absolute-move terms the same history reads 1.9, 2.9, 8.0, and 6.0 basis points across the four regimes, as Figure 4 draws year by year over the whole quarter century. The repricing of CPI days — and, per unit of surprise, its persistence through 2024–26 — is in the raw market data.

Figure 4. How big a day is CPI day? A consensus-free history, 1999–2026



Note: Annual mean absolute change in the 2-year Treasury yield on CPI-release days and on days with no major release, from FRED release calendars and daily yields; no consensus data enter. Source: Authors' calculations.

Employment days ran the opposite arc. They were the big days for most of the record — the 2-year multiple was 2.25 over 1999–2012 and 3.34 in the ZLB tail (the taper years) — then muted to 1.56 in the pre-COVID normal and 2.52 in the fight. In the normalization the employment-day multiple reached 5.29, with a mean absolute 2-year move of 8.5 basis points, larger than CPI days at their fight-era peak. The biggest scheduled mornings on the calendar in 2024–26 are jobs mornings.

Equities tell the same two stories without a single consensus number entering the calculation. The S&P 500's CPI-day multiple was below one in both pre-COVID regimes — 0.57 in 2016–19, a reminder that these are noisy ratios and that pre-COVID CPI mornings were, if anything, quiet — then 2.23 in the fight and 1.10 in the normalization — the equity repricing and the equity reversion both appear in raw variance. Its employment-day multiple went from 1.06 in the fight to 2.39 in the normalization — the labor rotation, in equities this time.

The multiple also answers the most natural objection: that everything simply became more volatile after COVID and per-print sensitivity merely scaled with the times. The multiple is a ratio. Baseline non-release-day variance did rise post-COVID, and release days rose by more. If sensitivity had scaled with general uncertainty, the multiples would have stayed flat. They more than doubled.

There is a sharper version of the objection. The whole calendar got bigger in the normalization — the any-release multiple on the 2-year rose from 1.31 in the fight to 2.01 — so measured against the average release day rather than against quiet days, CPI-day distinctiveness peaked in the fight at 1.96 times the average release day and narrowed to 1.37 in the normalization. Both readings are true, and the reconciliation is the rotation: CPI days remain nearly three times an ordinary day, and their edge over other releases narrowed because employment days surged past them, not because inflation days quieted.

Two cautions apply to Table 4. Days of FOMC announcements are not excluded from the baseline, which inflates the denominator and biases every multiple toward one; the reported numbers understate announcement-day distinctiveness. And one cell of Table 4 — the employment-day multiple for the 10-year breakeven, 3.49 in the normalization — is a variance ratio computed on a low-variance series and should not be read as contradicting the beta evidence of the next section, which it is too crude to overturn.⁴

Both legs of the design, then, deliver the same dating and the same verdict: the shift arrived in 2021–22; per unit of surprise the rates-side repricing persisted through August 2026, even as shrinking surprises returned CPI days most of the way to ordinary size; the equity response reverted; and the calendar’s weight shifted back toward employment. Whatever is

⁴ The breakeven variance multiple counts all movement on employment days, whatever its source, including the liquidity and risk-premium components of TIPS pricing that Section VI’s regression evidence conditions on the surprise itself.

wrong with the consensus feed, it is not driving the paper's facts.

VI. Real rates, not breakevens: what the decomposition says about anchoring

A CPI-day move in nominal yields can mean two very different things. It can be repriced inflation compensation — investors marking up the inflation they expect to be compensated for, the signature of an expectations anchor coming loose. Or it can be a repriced real policy path — investors marking up what the Federal Reserve will do about inflation, holding long-run inflation expectations fixed. TIPS yields and breakeven inflation rates split the nominal move into exactly these two components at the daily frequency. The pre-COVID benchmark comes from Beechey and Wright (2009), using intraday TIPS data: price news moved inflation compensation, modestly, while real-activity news moved real rates.

Table 5. The nominal response decomposed: real rates versus breakeven inflation

		Pre-COVID (2013-Feb 2020)	2021-23	2024-26
CPI (m/m)	10-year	0.04 (1.13) [26]	6.7** (3.4) [12]	2.31 (1.89) [10]
	10-year breakeven	1.16*** (0.37) [26]	1.97*** (0.60) [12]	1.01** (0.46) [10]
	10-year real	-1.11 (1.01) [26]	4.7 (3.5) [12]	1.31 (1.49) [10]
	5-year breakeven	0.95* (0.53) [26]	3.6*** (0.8) [12]	2.11*** (0.75) [10]
Core CPI (m/m)	10-year	-1.58 (1.23) [12]	5.2** (2.2) [14]	6.0*** (1.9) [10]
	10-year breakeven	1.99** (0.97) [12]	2.05*** (0.68) [14]	0.87 (0.65) [10]
	10-year real	-3.6*** (0.9) [12]	3.2 (2.3) [14]	5.2*** (1.4) [10]
	5-year breakeven	3.2* (1.9) [12]	3.5*** (0.9) [14]	2.04*** (0.57) [10]
Core PCE (m/m)	10-year	-2.01 (1.70) [11]	3.9*** (1.3) [7]	0.70 (0.76) [11]
	10-year breakeven	-1.39** (0.60) [11]	-0.72 (1.44) [7]	-0.37 (0.35) [11]
	10-year real	-0.62 (1.33) [11]	4.6*** (0.7) [7]	1.07* (0.57) [11]
	5-year breakeven	-1.68* (0.90) [11]	-0.91 (1.23) [7]	0.12 (0.37) [11]
PPI (m/m)	10-year	0.62 (0.83) [30]	3.6 (2.8) [10]	1.87*** (0.73) [13]
	10-year breakeven	0.73** (0.31) [30]	0.39 (0.92) [10]	0.61* (0.31) [13]
	10-year real	-0.11 (0.85) [30]	3.2 (2.2) [10]	1.26*** (0.47) [13]
	5-year breakeven	1.97*** (0.39) [30]	0.78 (1.04) [10]	0.88*** (0.25) [13]

Note: Cells as in Table 2. The 10-year nominal response is decomposed into the 10-year TIPS (real) response and the 10-year breakeven response; the 5-year breakeven is shown as the near end of the inflation-compensation curve. Breakevens are close-of-day constant-maturity series and carry liquidity and risk premia, discussed in the text. Source: Authors' calculations.


In our pre-COVID data — 2013–15, the years with consensus coverage for core CPI — the breakeven column shows the anchored-era pattern: a core CPI surprise moved the 10-year breakeven 2.0 basis points per standard deviation. The real-rate column shows something the Beechey-Wright benchmark does not predict: a negative response of 3.6 basis points, a growth-worry reading of hot inflation prints at the ZLB margin, and a deviation that matters when reading the pre-COVID nominal response of roughly zero. Before COVID, the two components of a CPI surprise nearly canceled: modest inflation compensation up, real rates

down, nominal yields unmoved.

The striking fact about the inflation fight is what did not amplify. At the daily frequency the breakeven response to core CPI in 2021–23 was 2.1 basis points per standard deviation — essentially unchanged from the pre-COVID 2.0 — while the real-rate response swung from minus 3.6 to plus 3.2. The nominal amplification everyone traded through was, on this decomposition, a real-rate phenomenon from the start.⁵

The normalization delivers the clean result, as shown in Figure 5: the 10-year response to core CPI in 2024–26 is 5.2 basis points per standard deviation in the real rate and 0.9 in the breakeven. A hot core CPI print now moves TIPS yields nearly one-for-one with nominal yields, and inflation compensation barely at all: the repricing is about the expected policy path, not the long-run inflation view.

Figure 5. Where the 10-year response to core CPI surprises comes from



⁵ Kroner (2025) finds larger responses of market-implied inflation expectations to CPI surprises in 2021–23 using intraday windows on inflation swap rates, an instrument he prefers to TIPS breakevens on liquidity grounds. Daily close-only breakevens smooth intraday round trips and carry liquidity premia, so the two findings measure different objects; the normalization-era verdict below does not depend on which is preferred for the fight years.

Note: OLS responses of the 10-year nominal yield, the 10-year TIPS yield, and the 10-year breakeven to a one-standard-deviation core CPI surprise, by regime; whiskers are 95 percent confidence intervals on HC1 standard errors. Source: Authors' calculations.

The term structure of inflation compensation says the same thing from another angle.

Five-year breakevens respond more than 10-year breakevens in every post-COVID regime — 3.5 against 2.1 in the fight, 2.0 against 0.9 in the normalization. Inflation news lives at the near end of the inflation curve and dies out by the long end, which is what anchoring means operationally. Compare Gürkaynak, Sack, and Swanson (2005): in the 1990s and early 2000s, far-forward nominal rates responded strongly to news they should have ignored if long-run expectations were anchored. By that classic diagnostic, anchoring improved across the decades in which CPI-day volatility exploded — the two facts coexist because the volatility is about the policy path, not the long-run level.

One piece of arithmetic disciplines the whole breakeven column. TIPS principal is indexed to the realized price level, so a core CPI surprise of one standard deviation — about 0.13 percentage point on the level, permanent even if the inflation rate itself carries no persistence — mechanically raises a 5-year breakeven by roughly 2.6 basis points and a 10-year by roughly 1.3, with no change in expected future inflation at all. The measured responses sit near — and in the normalization below — that mechanical floor, and the 5-year-greater-than-10-year gradient is the one-over-maturity scaling the indexation channel predicts. Netting the floor out, the expectations component of the breakeven response is approximately zero throughout — which strengthens the anchoring conclusion and cautions against reading even the pre-COVID 2.0 as economically meaningful inflation-compensation news. (Beechey and Wright's use of forward rates exists partly to strip this channel.)

The same pattern holds across releases — the thin cells are flagged in Table 5, and the ISM decomposition sits in the appendix grid. Core PCE responses in the post-COVID period are

real-rate-dominated. And the ISM index's fight-era surge ran almost entirely through real rates — 9.5 basis points per standard deviation real against minus 0.4 breakeven — which is what one would expect if activity news mattered in 2021–23 because of what it implied the Fed would do.

Two caveats apply. Breakevens are close-of-day constant-maturity approximations, and TIPS carry liquidity and risk premia that widen in stressed markets, so daily breakeven moves are an imperfect proxy for expected inflation, particularly in the fight years. And the decomposition is descriptive — it identifies which component moved, not a structural term-structure model of why. Neither caveat touches the normalization result, which is the one the argument rests on.

The inflation-expectations anchor held — on this evidence it never seriously slipped even in the fight — and the post-COVID sensitivity of yields to inflation news is a statement about the perceived reaction to inflation, not about inflation expectations coming loose. Section VIII builds on exactly this distinction.

VII. Equities and the sector cross-section

Equity responses need a different interpretive frame than yield responses. McQueen and Roley (1993) first showed that the equity response to macro news flips sign with the state of the economy, and Boyd, Hu, and Jagannathan (2005) supplied the frame: every macroeconomic print bundles news about corporate cash flows, news about the discount rate, and news about risk premia, and the equity response is the net of the three. The bundle is why equity responses flip sign with the state of the economy — good real news is good for stocks in a recession and bad for them in a boom — and why a flat equity response can hide two large offsetting channels. Before COVID, the S&P 500's response to a hot CPI print was almost exactly zero. The bundle netted out.

The inflation fight collapsed the bundle to its rate leg. A one-standard-deviation core CPI surprise cost the S&P 500 1.3 percent in 2021–23. The fitted response at September 13, 2022’s 2.3-standard-deviation surprise — about 2.9 percent — carries the sign and roughly two-thirds of that day’s 4.4 percent decline, with the remainder the bundle’s risk-premium leg on a morning that also repriced uncertainty about the Federal Reserve itself. When the binding risk in the economy is what the central bank will do about inflation, an inflation print is rate news.

The equity response then did something the rates response did not: it reverted. In the normalization the S&P 500’s response to core CPI is minus 0.1 percent per standard deviation — indistinguishable from zero and from its pre-COVID self. The fight-to-normalization interaction is 1.2 percentage points with p equal to 0.013, and the consensus-free record says the same thing on roughly seventy release days: the S&P 500’s CPI-day variance multiple fell from 2.23 to 1.10. The two legs together make the equity reversion the sharpest reversion fact in the paper. Once the policy path stopped being the binding risk, equities stopped trading the CPI print.

The post-COVID equity response is asymmetric in degree, though not in the way a sample with missing relief rallies would suggest. Splitting surprises by sign, the 2-year yield responds symmetrically to core CPI — 8.0 basis points per standard deviation for hot prints, 8.4 for soft — while equities lose 1.35 percent per standard deviation on hot prints and gain 0.56 on soft ones (Appendix A2). Hot prints hurt roughly two and a half times as much as soft prints helped. The two hand-restored mornings are what discipline this comparison: without August 10 and November 10, 2022 — among the cycle’s largest relief rallies — the soft-print response

appears negligible and the equity market looks like a one-way bet, which is precisely the error a non-random data gap manufactures.⁶

Payroll surprises moved equities little in every regime — 0.3 percent per standard deviation before COVID, 0.1 after. The bundle explains why: a strong payroll print is good news for earnings and bad news for rates, and the two legs offset. The employment-day variance multiple rising to 2.39 in the normalization while the payroll beta stays near zero is the bundle made visible — jobs mornings carry a lot of equity news, with a sign that depends on which leg dominates that morning. A release can matter enormously and still have a small average coefficient; the variance evidence complements the beta evidence rather than contradicting it.

The sector cross-section turns the aggregate response into a mechanism test, because the discount-rate channel makes a sharp prediction: sectors whose values sit in long-duration future cash flows, and sectors that trade as bond proxies, should lose the most when a print raises the expected policy path. Table 6 and Figure 6 report the core CPI responses by sector for the inflation fight, in percent per standard deviation on fourteen observations each: technology minus 1.78, consumer discretionary minus 1.76, communication services minus 1.46, real estate minus 1.34, materials minus 1.22, industrials minus 1.08, financials minus 1.04, energy minus 0.79, utilities minus 0.79, consumer staples minus 0.65, and health care minus 0.56. The gradient runs from the long-duration and cyclical-discretionary sectors, at two to three times the defensive sectors' losses, down to staples and health care.

⁶ Three robustness notes. Excluding core CPI mornings that share the calendar with another major release leaves the fight and normalization responses at 7.4 and 7.5 basis points per standard deviation. Measuring the normalization equity response on the price-index-only series (no total-return splice) gives the same minus 0.1. And the anomalous positive equity response to headline CPI in 2016–19 vanishes when the two largest surprises are dropped (the slope falls to essentially zero), confirming it as a small-cell artifact.


Table 6. Sector equity responses to surprises

		Pre-COVID (2013-Feb 2020)	2021-23	2024-26
CPI (m/m)	Comm. services	[n=1]	-2.17*** (0.66) [12]	[n=3]
	Cons. discretionary	0.04 (0.22) [26]	-2.24*** (0.85) [12]	[n=3]
	Cons. staples	-0.22 (0.17) [26]	-0.82* (0.42) [12]	[n=3]
	Energy	0.08 (0.18) [26]	-1.08*** (0.25) [12]	[n=3]
	Financials	0.18 (0.27) [26]	-1.57*** (0.55) [12]	[n=3]
	Health care	-0.06 (0.19) [26]	-0.87* (0.44) [12]	[n=3]
	Industrials	0.03 (0.18) [26]	-1.48*** (0.47) [12]	[n=3]
	Materials	0.03 (0.23) [26]	-1.72*** (0.51) [12]	[n=3]
	Real estate	-0.09 (0.15) [9]	-1.87** (0.90) [12]	[n=3]
	S&P 500	0.01 (0.19) [26]	-1.77*** (0.61) [12]	-0.18 (0.22) [10]
	Technology	0.05 (0.20) [26]	-2.53*** (0.84) [12]	[n=3]
	Utilities	-0.05 (0.18) [26]	-1.08 (0.67) [12]	[n=3]
Core CPI (m/m)	Comm. services	[n=0]	-1.46*** (0.53) [14]	[n=1]
	Cons. discretionary	-0.34 (0.28) [12]	-1.76*** (0.52) [14]	[n=1]
	Cons. staples	-0.22 (0.28) [12]	-0.65** (0.30) [14]	[n=1]
	Energy	-0.40 (0.55) [12]	-0.79*** (0.28) [14]	[n=1]
	Financials	-0.33 (0.30) [12]	-1.04** (0.44) [14]	[n=1]
	Health care	-0.35 (0.33) [12]	-0.56* (0.34) [14]	[n=1]
	Industrials	-0.23 (0.22) [12]	-1.08*** (0.34) [14]	[n=1]
	Materials	-0.36 (0.31) [12]	-1.22*** (0.40) [14]	[n=1]
	Real estate	[n=0]	-1.34** (0.56) [14]	[n=1]
	S&P 500	-0.32 (0.26) [12]	-1.27*** (0.43) [14]	-0.11 (0.18) [10]
	Technology	-0.33 (0.22) [12]	-1.78*** (0.58) [14]	[n=1]
	Utilities	-0.05 (0.28) [12]	-0.79** (0.40) [14]	[n=1]
Nonfarm payrolls	Comm. services	-0.08 (0.25) [8]	0.01 (0.11) [15]	[n=5]

	Pre-COVID (2013-Feb 2020)	2021-23	2024-26
Cons. discretionary	0.28 (0.24) [46]	0.16 (0.19) [15]	[n=5]
Cons. staples	-0.00 (0.17) [46]	-0.06 (0.10) [15]	[n=5]
Energy	0.56** (0.25) [46]	0.23 (0.20) [15]	[n=5]
Financials	0.67*** (0.23) [46]	0.27** (0.13) [15]	[n=5]
Health care	0.23 (0.21) [46]	0.07 (0.08) [15]	[n=5]
Industrials	0.42* (0.23) [46]	0.20* (0.12) [15]	[n=5]
Materials	0.25 (0.22) [46]	0.25** (0.13) [15]	[n=5]
Real estate	-0.01 (0.21) [20]	0.12 (0.13) [15]	[n=5]
S&P 500	0.30 (0.19) [46]	0.13 (0.10) [15]	0.15 (0.35) [17]
Technology	0.28 (0.21) [46]	0.15 (0.16) [15]	[n=5]
Utilities	-0.84*** (0.25) [46]	0.07 (0.12) [15]	[n=5]

Note: Cells as in Table 2, for the eleven Select Sector SPDR funds and the S&P 500. The CRSP daily file ends December 31, 2024, so the 2024–26 sector cells contain 2024 only and are printed for completeness, not interpretation. Source: Authors' calculations.

Figure 6. The cross-section of sector responses to core CPI surprises, 2021–23



Note: OLS responses per one-standard-deviation core CPI surprise; whiskers are 95 percent confidence intervals on HC1 standard errors. Gold bars mark the classic rate-sensitive sectors (utilities, real estate, financials); the red bar is the S&P 500. Source: Authors' calculations.

The ordering invites a duration reading, but it needs a sharper test before it can carry one, because technology and consumer discretionary are also the highest-market-beta sectors — on any morning the aggregate falls 1.3 percent, mechanical beta produces most of this ranking by itself. Adjusting each sector’s release-day return for its market beta, estimated on non-release days within the fight, answers the question directly: the abnormal responses are small and mostly indistinguishable from zero (Appendix A2). The CPI-day equity repricing was, first and foremost, an aggregate discount-rate shock that sectors inherited in proportion to their market betas. The residuals that survive the adjustment point the expected ways: financials outperform their beta on hot prints — plus 0.18 percent per standard deviation, the one significant abnormal response, the classic rate-margin channel — health care outperforms, real estate and technology tilt negative, and energy tilts positive once its beta drag is removed. The sector that sells the gasoline fell on gasoline-driven CPI mornings by no more than its market beta implies — and slightly less.

Pre-COVID payroll days show the same pattern with the roles reversed. On 2013–19 employment mornings, financials gained 0.67 percent per standard deviation of payroll surprise — the classic trade in which strong jobs mean higher rates and wider margins — while utilities, the bond proxy, lost 0.84. That gradient compressed in the fight, to 0.27 and 0.07 respectively: payrolls carried little rate news in 2021–23, so the rate-sensitivity map went quiet exactly when the CPI map lit up. Which release activates the cross-section is itself a regime variable.

The sector panel stops early: the CRSP daily file ends in December 2024 (Section III), so its normalization cells contain a single year, between one and five events per release — printed in Table 6 for completeness, leaned on nowhere. The aggregate equity series continues through August 2026 on the S&P 500 index level, which is what the reversion evidence above uses. A

proper sector-level reversion test needs intraday data past 2024, and Section X names it as the natural extension.

The cross-sectional map matters for interpretation because it parallels the map Bernanke and Kuttner (2005) and Ehrmann and Fratzscher (2004) drew for monetary-policy surprises — and because the beta adjustment locates the transmission: what a hot print delivers is an aggregate discount-rate shock, scaled across sectors by market beta, with a residual rate trade in financials. Markets in the fight treated a hot CPI print the way they treat a small surprise tightening — which is the perceived-rule story told in the cross-section.

Equities confirm both halves of the account — the repricing in the fight and the reversion in the normalization — and locate the mechanism in the discount rate. The next section scores the three stories on the evidence that can separate them.

VIII. Attention, perceived rule, or signal weighting: what the reversion test says

Three interpretations of the post-COVID repricing are on the table, and each carries a distinct prediction for 2024–26. The attention story (Kroner 2025) holds that sensitivity tracks the salience of inflation to investors; it predicts decay as inflation leaves the front page. The perceived-rule story (Bauer, Pflueger, and Sunderam 2024; Cline and Jia 2025) holds that sensitivity tracks how aggressively markets believe the Federal Reserve responds to the data; it predicts persistence until the perceived rule itself changes, with sensitivity rotating toward whichever side of the mandate is binding. The signal-weighting story (Hördahl, Kısacıkoglu, and Xia 2026) holds that sensitivity rises with monetary-policy uncertainty and falls with forecaster disagreement; it predicts sensitivity tracking the level of policy uncertainty.

The facts to be explained are now two-dimensional. The level of data-dependence stayed high: the any-release variance multiple rose to 2.01 in the normalization, payroll sensitivity reached its futures-era maximum, and per-unit sensitivity to core CPI held at 7.5 basis points per standard deviation with full reversion rejected on both legs. The composition rotated: the broader inflation complex gave back roughly three-fifths of its repricing (the composite fell from 6.9 to 2.8, significantly), the ISM was released from inflation duty, employment days became the calendar's biggest, and the equity CPI trade went quiet. And the decomposition says that whatever persisted is repricing of the real policy path, with inflation compensation pinned throughout. The perceived-rule story is the most parsimonious account of that two-part pattern: a market that stayed print-driven while redistributing weight across prints as the mandate's binding side shifted. The equity reversion is consistent corroboration, though under the bundle framing of Section VII it would follow from any story once the path stopped being the binding risk, so we do not count it as discriminating evidence.

The attention story has direct evidence behind it. Kroner's evidence that investor attention amplified the 2021–23 CPI response — news-coverage intensity, search behavior, and trading volume around CPI releases — is intraday, direct, and convincing for the peak of the episode. Our own fight-era equity multiple of 2.23 and its reversion to 1.10 fit an attention cycle riding on top of a repriced rule. The two mechanisms are complements operating at different frequencies: attention explains why 2022 CPI mornings were the biggest in the sample; the perceived rule explains why 2025 CPI mornings still dwarf 2019 ones.

What the evidence cannot do is reject an attention story in which salience tracks the mandate. Inflation never fully left the front page in 2024–26: core inflation was still above 3 percent for most of 2024, and the 2025 tariff episode put inflation back at the top of the news

cycle; the mid-2024 rise in the unemployment rate made jobs day front-page news at exactly the moment our payroll sensitivity re-emerges. An attention story in which the news cycle follows what the Federal Reserve cares about reproduces almost every row of the scorecard above, because the news cycle does follow what the Federal Reserve cares about. On this sample the two stories are close to observationally equivalent, with two partial exceptions that cut in opposite directions. The 2025 tariff episode kept inflation on the front page while the composite inflation response and the CPI-day multiple fell — salience without sensitivity, which fits a market tracking the mandate better than one tracking the news. Core CPI's unique persistence cuts the other way: if the Fed's rule weights the PCE index it targets, a market pricing the rule should not respond so much more durably to core CPI than to core PCE — unless the most famous print carries an attention premium of its own. What would separate them: attention proxies — news-coverage counts, search intensity — falling while rates sensitivity holds; a within-normalization split of the betas, which our cells are too thin to support; or the next easing cycle, when the mandate and the front page may finally point in different directions. The claim is that the perceived-rule story is the most parsimonious reading of the evidence, not that the alternatives are excluded.

The signal-weighting story is in the same position. Monetary-policy uncertainty remained elevated through 2024–26, so the persistence of inflation-release sensitivity is consistent with the Hördahl-Kısacıkoglu-Xia mechanism as well; a period of visibly low policy uncertainty with intact sensitivity would separate the stories, and this sample does not contain one.

Was the market right to trade the print?

One of the present authors is on record against exactly the behavior this paper documents. “One Good Inflation Report Isn’t Enough” (Furman 2022), written the day after a flat July 2022 CPI print set off one of the largest CPI-day rallies — one of the two hand-restored mornings in this sample — argued that one good report established little about underlying inflation and that the Federal Reserve should keep tightening aggressively rather than treat a single print as the turn. The implied warning — that a single print proves little, and that repricing policy on it gets ahead of the evidence — sits in obvious tension with a market moving 8.5 basis points per standard deviation of a single print.

The decomposition supplies the distinction that resolves it. A single print is a weak signal of underlying inflation — on that, the op-ed was and is correct, and the back-of-envelope below quantifies it. But when the central bank has committed to react to prints, the print is a nearly noiseless signal of policy. An investor on the morning of September 13, 2022 did not need to believe that August’s 0.6 percent core reading revealed much about underlying inflation; the investor needed only to believe that the FOMC would respond to it, which it did. Markets moving 8.5 basis points per standard deviation were not over-reading the inflation data. They were correctly reading a Fed that was reading the inflation data. Furman’s own later piece arguing that tightening should respond to the surprising data of early 2023 (Furman 2023) states the same reaction function from the prescriptive side: if surprising prints should move policy, then surprising prints should move markets.

A back-of-envelope calculation checks the magnitudes. A one-standard-deviation core CPI miss is about 0.13 percentage point on the monthly print, which moves the twelve-month inflation rate by about 0.1 percentage point mechanically — small, exactly as the noise argument

says. Under a reaction function that moves the expected policy path roughly one-for-one with revisions to the inflation outlook, and with the outlook revision to a persistent inflation process a multiple of that mechanical effect, the implied move in a 2-year yield that averages the expected path over its horizon is of the same order as the observed 7.5 to 8.5 basis points. The observed responses are not obviously excessive for that rule. The genuine question the 2022 op-ed raises is therefore about the rule, not the market: whether the Federal Reserve's own per-print reactivity over-weights noisy prints. That is a question about optimal policy under data-dependence, and it is beyond this paper's evidence — but the market's role in it is that of messenger.

Three observable developments would falsify the perceived-rule reading going forward. Sensitivity persisting after the Federal Reserve visibly stops being print-driven would break the link between the rule and the response. Breakevens rather than real rates carrying a renewed CPI response would say expectations, not policy, had become the margin. And the labor rotation failing to track which side of the mandate binds would decouple sensitivity from the rule's weights. All three can be tracked in this design as data accrue.

The sample's earliest years supply perspective. The payroll response of 5.1 basis points per standard deviation in the ZLB tail shows that the mechanism is old. Sensitivity concentrates wherever the policy trigger lives: in 2013 it lived in the exit timetable and payrolls; in 2022 it lived in the inflation fight and the CPI; in 2025 it lives in both sides of the mandate at once. Data-dependence long predates COVID; the COVID era moved it to inflation releases and increased its size.

The reversion test, run across rates, the decomposition, the activity releases, equities, and the raw calendar, reads the post-COVID repricing as a durable change in what markets believe the Federal Reserve responds to — amplified, briefly and measurably, by attention; not an

un-anchoring of inflation expectations; and not an artifact of a more volatile world.

IX. Implications

Implication 1: A data-dependent Federal Reserve outsources volatility to the statistical calendar

When policy credibly keys on a handful of prints, the variance of those prints' release days is a policy outcome. The Federal Reserve does not publish the CPI, but its reaction function determines whether a CPI morning is a 2.5-basis-point day or an 8-basis-point day, and 2021–26 demonstrates that the difference is chosen, not given. The announcement-premium literature says that concentrated event-day risk is priced — Savor and Wilson (2013) found that most of the equity premium accrues on scheduled announcement days, and Ai and Bansal (2018) showed such premia identify a departure from time-separable expected utility — preferences that price the resolution of uncertainty itself — so data-dependence carries a financial price, paid in risk premia, whether or not the central bank acknowledges setting it.

The practical corollary — if the perceived-rule reading of Section VIII is right — is that communication that flattens the perceived per-print rule is volatility policy: averaging language, total-inflation framing, explicit reaction-function talk, and guidance that distributes weight across many indicators all reduce the policy content of any single print, while meeting-by-meeting data-dependence concentrates it. The sample itself brackets the point: the averaging framework adopted in 2020 did not prevent the fight-era repricing once policy became print-driven, and the revised framework statement of August 2025 retired that language during the normalization — the perceived rule is set by what the Federal Reserve does about prints at least as much as by what its framework documents say. Under the rival attention reading, the volatility dial sits with

the news cycle rather than with communication at all — which is why Section VIII’s discriminating tests matter for policy and not only for interpretation. There are excellent reasons for data-dependence in an uncertain world, and this paper takes no position on the optimal rule. It establishes the price schedule: surprising the market has become expensive at precisely the frequency — twelve mornings a year — where the tools to manage it are communications rather than rate moves.

Implication 2: Announcement-day moves are now a live gauge of the perceived reaction function

Estimating what markets believe the Federal Reserve will do has traditionally required surveys of forecasters, as in Bauer, Pflueger, and Sunderam (2024), with the cadence of survey collection and the gap between what forecasters say and what traders do. The betas in this paper are computable from largely public data with a lag of months rather than quarters — though on roughly twelve observations a year per release, so the gauge is a complement to survey-based estimates, sharpest when read jointly with the consensus-free multiples, which accumulate evidence faster. In 2024–26 the market’s reweighting — labor re-emerging as the dominant release — was visible in prices throughout the normalization; dating it precisely against the communication record is a natural extension this design supports.

One interpretive discipline belongs with the gauge. An announcement beta is an equilibrium object, jointly determined by the perceived rule and by the signal content of the print — so a falling CPI beta can mean a flatter perceived rule or a noisier print, and Implication 4 documents that print quality genuinely deteriorated in 2025. The two implications are in deliberate tension: reading the gauge requires holding a view about which of its two determinants

moved, which is one more reason the consensus-free leg, the surprise-size record, and the beta belong in one dashboard rather than three.

Implication 3: The two-sided mandate is visible in market prices again — watch the rotation, not the level

Within the paper’s account, the inflation fight’s single-issue calendar was the anomaly, and its unwinding is the normalization working. Employment days carrying the largest variance multiples we measure, the payroll beta at its sample maximum, the whole employment report — payrolls and the unemployment rate jointly — priced rather than one headline number, and the ISM released from inflation duty: each is the market pricing a central bank with two live objectives rather than one emergency.

For forecasters and commentators the operational lesson is that “CPI day is the biggest day” is already stale, and the accurate statement is both halves at once: CPI-day rates sensitivity remains at fight-era levels, and jobs day has overtaken CPI day as the calendar’s center of gravity. The level measures how strongly markets believe the Fed responds to data; the rotation measures which data. A forecaster needs both.

Implication 4: First-print quality is financial-stability infrastructure

Markets trade the first print within the hour, and this paper measures how hard. The first print is also the least accurate estimate the statistical agencies will ever publish for that reference period: post-COVID payroll first prints in particular have carried larger revisions amid falling survey response rates, a data-quality deterioration that Hördahl, Kısacıkoglu, and Xia (2026) argue explains part of the employment report's faded yield response in the fight years. The combination — maximal market weight on minimal-precision estimates — means noise in the print becomes repriced policy expectations within minutes and unwinds only as revisions dribble out. The sample contains the stress test: in 2025 the BLS curtailed CPI price collection in several cities for staffing reasons, the commissioner was dismissed after a payroll report carrying large downward revisions, and the fall shutdown delayed the September CPI and left the October report unpublished. A market that trades the print within the hour now faces prints that arrive late, thinner, or not at all — while pricing them, on this paper's estimates, more heavily than at any time on record.

The implication is that funding and methodology at the BLS and the Bureau of Economic Analysis (BEA) — survey response rates, revision transparency, the integrity of the release process — now carry direct asset-price externalities measured in basis points per release. A statistical agency in a data-dependent policy regime is, functionally, market infrastructure, and ought to be resourced like it.

X. Conclusion

This paper measured how financial markets respond to macroeconomic data surprises across the whole transmission chain — the expected federal funds rate, the yield curve, its real and breakeven components, the equity aggregate, and eleven sectors — on one consistent metric, from 2013 through August 2026, with every consensus-based estimate paired against a consensus-free counterpart. Three findings organize the evidence. Inflation news was repriced after COVID — from essentially nothing to 8.5 basis points per standard deviation of core CPI on the 2-year yield — and the repricing has come only partway back: full reversion is rejected on both legs of the design, the broader inflation complex has returned roughly three-fifths of the move, and per-unit sensitivity to core CPI is the one measure that shows no decline. The repricing ran through real rates rather than inflation compensation, so it measures the perceived reaction function of the Federal Reserve, not an un-anchoring of inflation expectations. And the equity response, which was the fight’s most visible symptom, reverted as the policy path stopped being the binding risk, while the calendar’s center of gravity rotated back to the employment report.

The verdict on 2024–26 is provisional by construction. It rests on two and a half years of data, with consensus cells as small as ten observations for the monthly releases, and what that supports is the finding as Section IV states it: full reversion rejected for core CPI on both legs of the design; partial reversion established for the broader inflation complex; no detectable decline for core CPI itself, inside wide bands. The design is built to be re-run. Intraday event windows from exchange data would sharpen every estimate and extend the sector panel past 2024; a licensed consensus history (Bloomberg, Action Economics) would fill the survey record’s gap years and retire the ZLB benchmark caveat outright; and the next full easing cycle — one driven

by a deteriorating labor market rather than by disinflation alone — will supply the out-of-sample test Section VIII describes.

“Why do markets overreact to data?” turns out to be the wrong question — the moves are about the right size for a market pricing a print-driven policy rule. The right question — what rule do markets believe the Federal Reserve is running? — has a measurable answer, and it is updated twelve mornings a year at 8:31.

Appendix

A1. The acute window

March–December 2020 is excluded from every pooled estimate in the paper, and Appendix Table A1 shows what inclusion would mean. Measured responses in the acute window are approximately zero across assets — not because markets ignored the data, but because consensus forecasts missed by amounts (millions of jobless claims) that stretch the surprise axis until any slope degenerates. The acute window tells us about the consensus process in a crisis, not about market sensitivity, and it is reported here rather than mixed into any regime.

Table A1. The acute COVID window (March–December 2020)

		2020*
CPI (m/m)	10-year	[n=5]
	2-year	[n=5]
	S&P 500 (%)	[n=5]
Initial jobless claims	10-year	-0.01 (0.01) [42]
	2-year	-0.01 (0.01) [42]
	S&P 500 (%)	0.02** (0.01) [42]
Nonfarm payrolls	10-year	0.09** (0.03) [10]
	2-year	0.04** (0.02) [10]
	S&P 500 (%)	0.03*** (0.01) [10]

Note: Cells as in Table 2. Estimates are shown for completeness; the text explains why they are uninformative about market sensitivity. Source: Authors’ calculations.

A2. Asymmetry and same-morning joint regressions

Splitting post-COVID core CPI surprises by sign, the 2-year yield response is symmetric — 8.0 basis points per standard deviation for hot surprises, 8.4 for soft — while the S&P 500 loses 1.35 percent per standard deviation on hot surprises and gains 0.56 on soft ones. On CPI mornings core and headline surprises identify separate coefficients in a joint regression only in the pooled post-COVID sample (seventeen common observations), where the 2-year loads 4.6 on core and 2.4 on headline; on employment mornings, payrolls and the unemployment-rate surprise both enter post-COVID (2.8 and minus 4.9) while the unemployment coefficient pre-COVID is nil (0.2 on forty-three observations). The market-beta-adjusted sector responses referenced in Section VII — each sector’s release-day return net of its market beta estimated on non-release fight days, regressed on the surprise — are in the replication package (output/estimates_sectors_adjusted.csv), as are the wild-bootstrap p-values and leave-one-out ranges for every headline cell, the fight-to-normalization interaction for all ten releases (the family within which the ISM and core PCE reversions sit), and the missingness audit comparing market moves on the CPI mornings the consensus feed covers against the mornings it misses (output/missingness_audit.csv).

A3. Construction detail

Hand-repaired days: the consensus and first-print values for August 10 and November 10, 2022 come from contemporaneous press accounts of the Dow Jones and Bloomberg surveys (CNBC, August 10, 2022; Fortune, November 10, 2022), with rows and source citations in `data/raw/manual_consensus_repairs.jsonl`. Re-anchoring: each feed row's date is replaced by the nearest true release date within a window of three days before to one day after the feed's bucket date, from the FRED release calendar for the release's agency series (CPI, PPI, employment situation, GDP, personal income and outlays, retail sales, weekly claims); rows with no calendar match are dropped and counted. ISM, which has no FRED calendar, is dated to the first business day of the month. Month-over-month/year-over-year assignment: each same-named row is assigned to the concept whose revised-series value for the reference month is nearest its released actual, with a 0.5 tolerance and the reference month taken as one month before the release month (two for delayed prints); rows matching neither concept are dropped and counted. Futures: the expected-rate change at horizon h months is the day-over-day change in 100 minus the settlement price of the contract for calendar month t plus h , computed within the same contract. Surprise standardization uses the full matched sample excluding March–December 2020. Releases on non-trading days are matched to the next trading day.

A4. The full grid

The complete estimates — all ten releases, all ten asset outcomes, all periods, with standard errors and cell counts — are in the replication package (`output/estimates.csv`), along with the sector grid (`output/estimates_sectors.csv`) and the variance spine (`output/variance_spine.csv`).

References

- Ai, Hengjie, and Ravi Bansal. 2018. "Risk Preferences and the Macroeconomic Announcement Premium." *Econometrica* 86 (4): 1383–1430.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega. 2003. "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange." *American Economic Review* 93 (1): 38–62.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega. 2007. "Real-Time Price Discovery in Global Stock, Bond and Foreign Exchange Markets." *Journal of International Economics* 73 (2): 251–277.
- Balduzzi, Pierluigi, Edwin J. Elton, and T. Clifton Green. 2001. "Economic News and Bond Prices: Evidence from the U.S. Treasury Market." *Journal of Financial and Quantitative Analysis* 36 (4): 523–543.
- Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam. 2024. "Perceptions About Monetary Policy." *Quarterly Journal of Economics* 139 (4): 2227–2278.
- Beechey, Meredith J., and Jonathan H. Wright. 2009. "The High-Frequency Impact of News on Long-Term Yields and Forward Rates: Is It Real?" *Journal of Monetary Economics* 56 (4): 535–544.
- Bernanke, Ben S., and Kenneth N. Kuttner. 2005. "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *Journal of Finance* 60 (3): 1221–1257.
- Boyd, John H., Jian Hu, and Ravi Jagannathan. 2005. "The Stock Market's Reaction to Unemployment News: Why Bad News Is Usually Good for Stocks." *Journal of Finance* 60 (2): 649–672.
- Campbell, Sean D., and Steven A. Sharpe. 2009. "Anchoring Bias in Consensus Forecasts and Its Effect on Market Prices." *Journal of Financial and Quantitative Analysis* 44 (2): 369–390.
- Cline, Alexander, and Chengcheng Jia. 2025. "Has the Market's Perception of the FOMC's Reaction Function Changed since the Onset of the COVID-19 Pandemic?" Federal Reserve Bank of Cleveland, Economic Commentary 2025-12.
- Ehrmann, Michael, and Marcel Fratzscher. 2004. "Taking Stock: Monetary Policy Transmission to Equity Markets." *Journal of Money, Credit and Banking* 36 (4): 719–737.
- Fleming, Michael J., and Eli M. Remolona. 1997. "What Moves the Bond Market?" Federal Reserve Bank of New York Economic Policy Review 3 (4): 31–50.
- Furman, Jason. 2022. "One Good Inflation Report Isn't Enough." *Wall Street Journal*, August 11.
- Furman, Jason. 2023. "To Fight Inflation, Fed Tightening Should Go Faster and Further." *Wall Street Journal*, March 3.

- Gil de Rubio Cruz, Antonio, Emilio Osambela, Berardino Palazzo, Francisco Palomino, and Gustavo Suarez. 2023. “Inflation Surprises and Equity Returns.” Working paper, Federal Reserve Board.
- Gilbert, Thomas, Chiara Scotti, Georg Strasser, and Clara Vega. 2017. “Is the Intrinsic Value of a Macroeconomic News Announcement Related to Its Asset Price Impact?” *Journal of Monetary Economics* 92: 78–95.
- Goldberg, Linda S., and Christian Grisse. 2013. “Time Variation in Asset Price Responses to Macro Announcements.” Staff Report 626. Federal Reserve Bank of New York.
- Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. 2005. “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models.” *American Economic Review* 95 (1): 425–436.
- Hördahl, Peter, Burçin Kısacıkoglu, and Fan Dora Xia. 2026. “Bond Yield Responses to Macro News: The Role of Macro Forecast Disagreement and Monetary Policy Uncertainty.” BIS Working Paper 1361. Bank for International Settlements.
- Hu, Grace Xing, Jun Pan, Jiang Wang, and Haoxiang Zhu. 2022. “Premium for Heightened Uncertainty: Explaining Pre-Announcement Market Returns.” *Journal of Financial Economics* 145 (3): 909–936.
- Kroner, T. Niklas. 2025. “How Markets Process Macro News: The Importance of Investor Attention.” Finance and Economics Discussion Series 2025-022. Board of Governors of the Federal Reserve System.
- Lucca, David O., and Emanuel Moench. 2015. “The Pre-FOMC Announcement Drift.” *Journal of Finance* 70 (1): 329–371.
- McQueen, Grant, and V. Vance Roley. 1993. “Stock Prices, News, and Business Conditions.” *Review of Financial Studies* 6 (3): 683–707.
- Piazzesi, Monika, and Eric T. Swanson. 2008. “Futures Prices as Risk-Adjusted Forecasts of Monetary Policy.” *Journal of Monetary Economics* 55 (4): 677–691.
- Rigobon, Roberto, and Brian Sack. 2008. “Noisy Macroeconomic Announcements, Monetary Policy, and Asset Prices.” In *Asset Prices and Monetary Policy*, edited by John Y. Campbell, 335–370. Chicago: University of Chicago Press.
- Savor, Pavel, and Mungo Wilson. 2013. “How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements.” *Journal of Financial and Quantitative Analysis* 48 (2): 343–375.
- Swanson, Eric T., and John C. Williams. 2014. “Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates.” *American Economic Review* 104 (10): 3154–3185.