Central-Bank News and Market Surprises

Futures, high-frequency event studies, and communication signals

Summer 2026

Computational Finance / FinTech Summer School

Day 7 | expectations and empirical measurement

Day 6 used contracts to reshape risk; Day 7 reads expectations from contracts

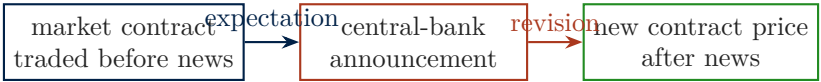

> A market surprise is measured from the change in an expectation-sensitive price around a known release time.

Expectation-sensitive price = a traded price whose value depends materially on beliefs about future rates or economic conditions.

29 July 2026: the target stayed unchanged, but short-rate expectations fell

Quantity	30-minute change	Interpretation
Target range	0 bp	maintained at 3.50–3.75%
MP1 surprise	−6.5 bp	current-meeting rate news was dovish
MP2 surprise	−6.5 bp	next-meeting expectation also fell
2-year Treasury yield	−4.5 bp	short-horizon discount rates declined
S&P 500	+0.14%	equity prices rose in the statement window

Federal Reserve; FRBSF USMPD, statement window 13:50–14:20 ET, updated 3 Aug 2026.

bp = basis point = 0.01 percentage point. MP1 = current-meeting target surprise; MP2 = expected-rate surprise at the next scheduled meeting.

Action, surprise, and reaction answer different questions

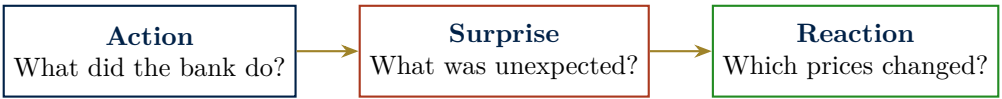

- The action is observed in the policy statement.
- The surprise requires a pre-announcement expectation measure.
- The reaction is an asset-price change measured over a declared event window.

A zero action does not imply a zero surprise; a zero surprise does not imply that every asset has a zero reaction.

Roadmap: six questions connect contracts, timestamps, prices, and words

1. **Expect:** what did the market believe before the announcement?
2. **Futures:** how does a contract quote reveal that belief?
3. **Surprise:** how do we isolate the unexpected component?
4. **Window:** what does a high-frequency event study identify?
5. **Words:** what does communication add beyond the current action?
6. **Text:** how can wording become an auditable measurement?

The sequence moves from the expectation object to measurement, identification, and interpretation.

By the end, you should be able to build and audit an announcement study

- Convert a 30-Day Federal Funds futures price into an implied monthly-average rate.
- Recover an expected post-meeting rate from calendar weights.
- Calculate and sign-check a current-meeting target surprise.
- Choose an event window and interpret yield changes, returns, and regression coefficients.
- Separate current-action news from news about the expected policy path.
- Construct a transparent text score and state what it does not identify.

Factor extraction, intraday microstructure, and structural monetary-policy identification appear only as linked extensions.

PART 1 OF 6

Expectations before news

What did the market believe before the announcement?

News is measured relative to a baseline. Define the belief and the horizon before using a price.

A policy action is news only relative to the expected action

$$\underbrace{\text{policy news}}_{\text{unexpected component}} = \underbrace{\text{announced action}}_{\text{observed after release}} - \underbrace{\mathbb{E}_{t^-}[\text{action}]}_{\text{belief just before release}} .$$

- A 25 bp cut can be no news if a 25 bp cut was fully expected.
- An unchanged target can be hawkish if markets expected a cut.
- The same unchanged target can be dovish if markets expected a rise.

t^- = instant just before the announcement; $\mathbb{E}_{t^-}[\cdot]$ = expectation using information available at that instant.

The event clock separates the baseline, the release, and the response

ET = U.S. Eastern Time. A scheduled FOMC statement is typically released at 14:00 ET; the press conference begins later and minutes are released on a different date.

Federal Reserve FOMC calendar and FRBSF USMPD window definitions.

Surveys, prices, and text observe different parts of expectations

Route	Direct object	Main limitation
Surveys	stated belief or forecast	infrequent, population-specific, sometimes stale
Market prices	traded valuation before and after news	risk premia, contract design, liquidity
Text	observable content of communication	measurement depends on document and method

> Day 7 uses market prices to measure belief revision and text to describe part of the information that caused it.

A method is informative only after the expectation object, population, horizon, and timing are declared.

Expected future overnight rates affect today's financial conditions

$$i_t^{(n)} \approx \frac{1}{n} \sum_{j=0}^{n-1} \mathbb{E}_t[i_{t+j}^{\text{overnight}}] + \tau_t^{(n)}.$$

- A central bank controls an overnight policy instrument most directly.
- Longer yields depend on expected future overnight rates across many dates.
- Term premia and other wedges prevent a pure-expectations interpretation.

$i_t^{(n)}$ = annualised yield over horizon n; $i^{\text{overnight}}$ = overnight rate; $\tau_t^{(n)}$ = term premium and other pricing wedges.

Appendix A8 · Pricing measure

The target range, the effective rate, and the futures rate are distinct

Object	Set or measured by	Role
FOMC target range	policy decision	desired range for the federal funds rate
Effective federal funds rate	New York Fed transaction data	realised overnight unsecured rate
Futures-implied rate	30-Day Federal Funds futures price	expected monthly average of the effective rate

Federal Reserve; Federal Reserve Bank of New York; CME contract specifications.

FOMC = Federal Open Market Committee. EFFR = effective federal funds rate, a volume-weighted median of eligible overnight federal funds transactions.

Appendix A2 · EFFR median

Different instruments reveal different horizons and risk exposures

- Futures and overnight indexed swaps are closest to the expected policy path.
- Treasury yields add term premia; equities and exchange rates also reflect growth, risk, and discount-rate news.

OIS = overnight indexed swap; FX = foreign exchange. Instrument choice must match the policy horizon being measured.

A market price is an expectation with wedges, not a survey answer

$$\text{market-implied object} = \text{expected payoff under pricing weights} + \text{contract and market wedges.}$$

- Risk premia can make pricing weights differ from real-world probabilities.
- Bid–ask spreads and thin trading add measurement noise.
- Settlement conventions determine exactly which rate or average is priced.
- Regulation, collateral, and balance-sheet costs can affect relative prices.

Pricing weights = probabilities adjusted for how states are valued by the market; they need not equal frequencies that investors expect to observe.

Appendix A8 · Q versus P

Before 29 July, the relevant baseline was the expected post-meeting rate

1. Observe the current target range and realised EFFR before the meeting.
2. Read a meeting-sensitive futures contract immediately before 14:00 ET.
3. Translate its monthly-average rate into an expected post-meeting rate.
4. Compare that expectation with the rate implied immediately after 14:00.

> The baseline is not “what analysts said last week”; it is the price-implied expectation at the chosen pre-event timestamp.

The exact contract and timestamp must be recorded so another researcher can reproduce the baseline.

Checkpoint: define four items before opening the dataset

1. **Object**: current target, next-meeting rate, full policy path, inflation, or another expectation?
2. **Horizon**: overnight, meeting month, one year, or longer?
3. **Population**: households, professional forecasters, or market participants?
4. **Timestamp**: survey date, pre-statement price, post-statement price, or full-event close?

> A precise expectation question determines the correct instrument and prevents interpretation after the result is known.

Read the futures contract

How does a quote reveal a policy expectation?

Start with settlement. The contract prices a monthly average, so the calendar enters every inference.

The contract settles on the monthly average effective federal funds rate

$$\bar{r}_m = \frac{1}{N_m} \sum_{j=1}^{N_m} r_{m,j}, \qquad \text{final settlement price} = 100 - \bar{r}_m.$$

- Each $r_{m,j}$ is the published EFFR for calendar day j in month m.
- The same rate is used for weekends and holidays as specified by the contract convention.
- The object is a monthly average, not the rate on the FOMC meeting day.

CME, 30-Day Federal Funds futures contract specifications.

N_m = number of calendar days in month m; $\bar{r}_m$ = monthly-average EFFR in annual percentage points.

Appendix A3 · Settlement identity

The quote convention subtracts the implied rate from 100

$$F_t = 100 - \mathbb{E}_t^Q[\bar{r}_m] \qquad \Longleftrightarrow \qquad \mathbb{E}_t^Q[\bar{r}_m] = 100 - F_t.$$

Futures price F_t	Implied monthly rate	Reading
94.50	5.50%	lower price, higher rate
95.00	5.00%	benchmark
95.50	4.50%	higher price, lower rate

Q = market pricing measure. The implied rate is annualised but applied to the contract month's daily average for settlement.

Python P2 · Price to rate

A futures-price rise means the expected rate has fallen

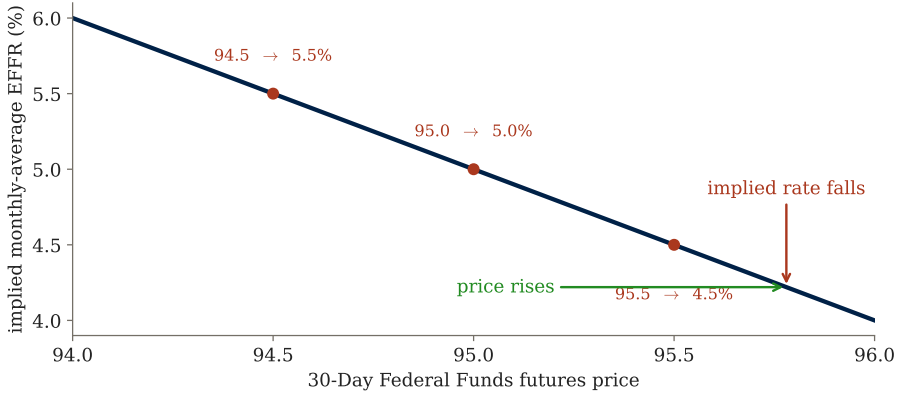

$\Delta \mathbb{E}^{Q}[\bar{r}_m] = -\Delta F$: the minus sign is essential when translating an intraday futures-price movement into a rate surprise.

Appendix A6 · Sign audit Python P2 · Price to rate

A meeting-month contract mixes rates from before and after the decision

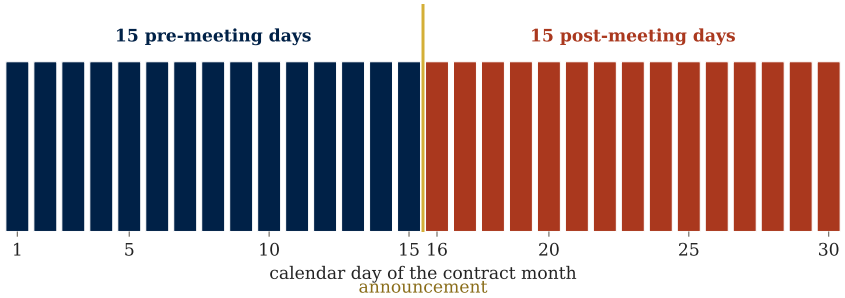

$$\mathbb{E}_t[\bar{r}_m] = \frac{d-1}{N}\, r_{\text{pre}} + \frac{N-d+1}{N}\, \mathbb{E}_t[r_{\text{post}}].$$

d = announcement day; r_{pre} = average rate before the decision; r_{post} = rate expected from day d onward.

Appendix A4 · Weight derivation

Calendar weights determine how much of the month can respond

Meeting day d	Pre-meeting weight	Post-meeting weight	Consequence
2	1/30	29/30	contract is highly meeting-sensitive
16	15/30	15/30	revision is diluted by one half
26	25/30	5/30	small contract move can imply large surprise
30	29/30	1/30	current-month contract becomes fragile

> Meeting timing is part of the financial instrument, not an administrative detail.

Weights assume a 30-day month and implementation from day d; applied work must check the exact effective date.

Before the meeting, most pre-meeting daily rates are already observed

- Published EFFR observations before the event are realised data, not forecasts.
- The remaining pre-meeting day can usually be approximated using the current operating environment.
- The unknown component of a meeting-month contract is concentrated in post-meeting days.

$$\underbrace{N\,\mathbb{E}_t[\bar{r}_m]}_{\text{implied monthly sum}} - \underbrace{(d-1)r_{\text{pre}}}_{\text{known pre-meeting contribution}} = \underbrace{(N-d+1)\mathbb{E}_t[r_{\text{post}}]}_{\text{unknown post-meeting contribution}} .$$

The decomposition is accounting. The economic interpretation enters when the post-meeting rate is mapped to possible FOMC outcomes.

Solve the monthly average for the expected post-meeting rate

$$\boxed{\mathbb{E}_t[r_{\text{post}}] = \frac{N\,\mathbb{E}_t[\bar{r}_m] - (d-1)r_{\text{pre}}}{N-d+1}}$$

1. Convert futures price into the implied monthly-average rate.
2. Remove the contribution of days before the meeting.
3. Divide the remaining rate sum by the number of post-meeting days.

Every rate in the formula uses the same annual percentage-point convention; the calendar weights are dimensionless.

Appendix A4 · Derivation Python P3 · Post-meeting rate

Worked example: a mid-month contract implies a 5.35% post-meeting rate

Input	Value	Meaning
Futures price	94.775	monthly rate = 5.225%
Month length N	30	calendar days
Announcement day d	16	15 days before, 15 after
Pre-meeting rate	5.10%	realised/near-known rate

$$\mathbb{E}_t[r_{\text{post}}] = \frac{30(5.225) - 15(5.10)}{15} = 5.35\%.$$

Illustrative inputs. Substituting 5.35% into the weighted average reproduces 5.225%.

Python P3 · Verify the weights

Discrete policy outcomes turn the expected rate into an implied probability

Suppose the post-meeting rate can be only 5.10% or 5.35%:

$$5.30 = p(5.35) + (1-p)(5.10) \qquad \Longrightarrow \qquad p = \frac{5.30 - 5.10}{5.35 - 5.10} = 0.80.$$

- The expected rate is a probability-weighted average of possible outcomes.
- The two-outcome mapping is exact only when those outcomes span the relevant state space.
- More possible outcomes require more prices, restrictions, or a distributional model.

Here p is the pricing-measure probability of the 25 bp higher outcome under the stated two-state assumption.

A market-implied probability is neither a promise nor a pure belief

- **Model dependence**: the answer changes if additional policy outcomes are allowed.
- **Pricing weights**: investors may value states differently, so Q probabilities can differ from real-world P probabilities.
- **Contract basis**: settlement is tied to the monthly average EFFR, not mechanically to the target midpoint.
- **Market conditions**: liquidity and bid–ask spreads limit spurious precision.

> Report the price, assumptions, timestamp, and outcome grid with every probability.

P = physical or real-world probability measure; Q = pricing measure used by market prices.

Appendix A8 · Q versus P

Late-month meetings can make the current contract an unstable measuring instrument

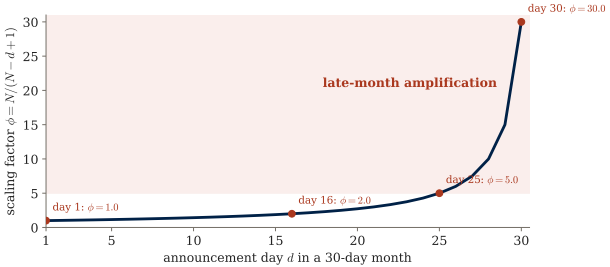

- Fewer post-meeting days mean less weight on the decision.
- Scaling amplifies both the signal and any quote or microstructure error.
- Researchers may combine adjacent contracts when the meeting is very late in the month.

$\phi = N/(N-d+1)$ is the scaling factor. Late-month corrections depend on adjacent contracts and the exact schedule.

Appendix A7 · Late-month correction

Checkpoint: keep four transformations in the correct order

> The quote is observed; each later object requires an additional convention or assumption.

Do not call $100 - F_t$ the meeting outcome before removing the pre-meeting days.

PART 3 OF 6

Measure the surprise

How much of the announcement was unexpected?

Use the price revision, correct for the contract calendar, and audit the sign before interpretation.

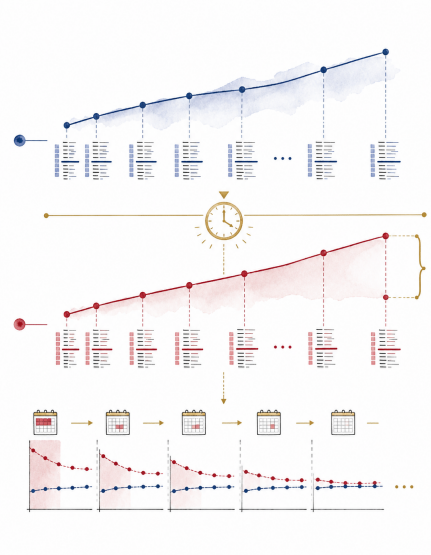

A surprise is a revision across the release, not the level of a rate

$$\text{surprise at } t = \underbrace{\mathbb{E}_{t^+}[r_{\text{post}}]}_{\text{after the release}} - \underbrace{\mathbb{E}_{t^-}[r_{\text{post}}]}_{\text{before the release}} .$$

- The before and after prices refer to the same contract.
- The event window should be short enough that the announcement dominates new information.
- The unit is a rate change, usually reported in basis points.

t^+ = instant just after the release. A positive rate surprise means the expected policy rate rose relative to the pre-release baseline.

The Kuttner measure scales the monthly-rate revision to the meeting

$$TS_t = \frac{N}{N-d+1}\,\Delta\widetilde{f}_t = -\frac{N}{N-d+1}\,\Delta F_t$$

- $\widetilde{f}_t = 100 - F_t$ is the contract-implied monthly-average rate.
- Δ means post-announcement value minus pre-announcement value.
- The scaling factor removes dilution from days before the meeting.

Kuttner (2001), adapted to an intraday before/after futures-price window.

TS_t = current-meeting target surprise; N = days in month; d = announcement day; F_t = 30-Day Federal Funds futures price.

Appendix A5 · Derivation Python P4 · Calculate surprise

Sign audit: a futures-price fall is a positive rate surprise

Futures price change	Implied-rate change	Surprise interpretation
$\Delta F < 0$	$-\Delta F > 0$	policy news is more hawkish than expected
$\Delta F = 0$	$-\Delta F = 0$	no measured current-contract rate news
$\Delta F > 0$	$-\Delta F < 0$	policy news is more dovish than expected

> Never attach "hawkish" or "dovish" to the raw futures-price sign; convert price into a rate first.

Hawkish = news associated with a higher expected policy-rate path; dovish = news associated with a lower expected path.

Appendix A6 · Unit and sign audit

Scaling converts a diluted monthly move into the meeting-specific move

Meeting day	Post weight	Monthly rate move	Meeting surprise
1	30/30	+5 bp	+5 bp
16	15/30	+5 bp	+10 bp
26	5/30	+5 bp	+30 bp

monthly revision = post-meeting weight × meeting-specific revision.

The table uses a 30-day month. A fixed 5 bp monthly revision carries very different information as the post-meeting weight changes.

Worked example: a 5 bp monthly revision becomes a 10 bp surprise

Input	Value	Transformation
Price before	94.775	implied rate 5.225%
Price after	94.725	implied rate 5.275%
Price change ΔF	-0.050	monthly rate change $+5$ bp
Calendar	$N = 30, d = 16$	scale $30/15 = 2$

$$TS = -\frac{30}{15}(-0.050) = +0.100 \text{ percentage point} = +10 \text{ bp}.$$

Price quotes and percentage-point rates share the numerical scale here; multiply the surprise by 100 to report bp.

Python P4 · Unit-tested function

The same 25 bp action can contain very different news

Pre-meeting expectation	Announced action	Surprise	Market description
0 bp change	+25 bp	+25 bp	hawkish surprise
+25 bp change	+25 bp	0 bp	fully anticipated
+50 bp change	+25 bp	−25 bp	dovish relative to expectation

> Regressing asset returns on the action confounds expected policy with unexpected policy.

Signs describe the surprise relative to the baseline, not whether the announced policy rate is high or low in an absolute sense.

An unchanged target can still create a large surprise

- The policy-rate action is zero.
- The surprise can be negative if the pre-release expectation was higher.
- Wording or dissents can also change expectations for later meetings.

This logic explains why the 29 July 2026 target decision and its measured MP1 surprise have different signs and magnitudes.

Recent events separate current-meeting news from next-meeting news

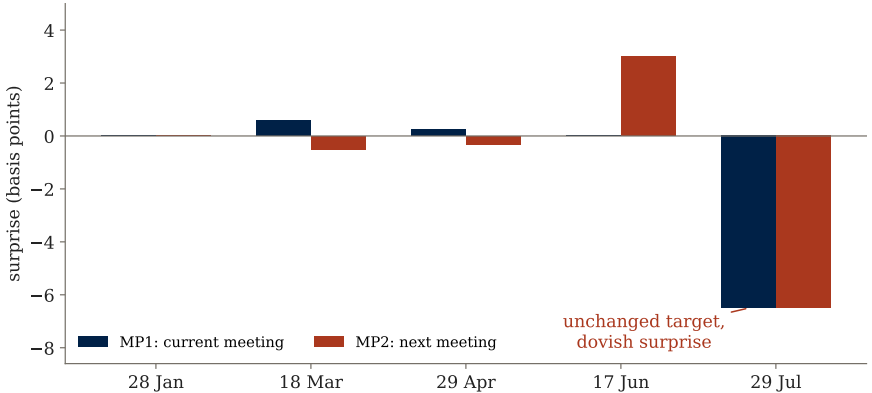

FRBSF USMPD, 30-minute statement windows; values converted from percentage points to bp.

The 17 June observation has zero MP1 but positive MP2; 29 July has negative MP1 and MP2. One contract cannot summarise the entire expected path.

Python P5 · Reproduce chart

MP1 answers a deliberately narrow current-meeting question

- MP1 measures the announcement-window revision attributed to the current meeting.
- It is designed to remove the meeting-month averaging structure.
- It does not by itself capture guidance about several future meetings.
- It does not reveal why the central bank changed the outlook.
- It does not prove a structural monetary-policy shock.

> A good measure can be narrow: precision about one object is better than an ambiguous "policy shock."

MP1 and MP2 are USMPD variable names; later-horizon futures and OIS changes carry additional path information.

Measurement errors enter through contracts, timing, and market conditions

- **Contract mismatch**: the chosen contract may have little post-meeting weight.
- **Implementation timing**: the effective rate may not move exactly when the simplified calendar assumes.
- **Liquidity**: stale quotes or wide spreads can dominate a small intraday revision.
- **Risk premia**: futures rates are pricing-measure objects, especially at longer horizons.
- **Multiple news**: the statement can change both policy expectations and the economic outlook.

These qualifications motivate reporting multiple contracts, exact windows, and robustness checks rather than a single unexplained series.

Checkpoint: the action is observed; the surprise is inferred

1. Record the released policy action and timestamp.
2. Read the same futures contract immediately before and after the release.
3. Convert the price change into an implied-rate change with the correct sign.
4. Apply the meeting-month scaling factor and convert percentage points to bp.
5. State the object measured and the information it omits.

> The surprise is a constructed measurement with an audit trail, not a field copied from the policy statement.

Isolate the event

What does a high-frequency event study identify?

Measure the surprise and the asset response over the same tightly controlled information interval.

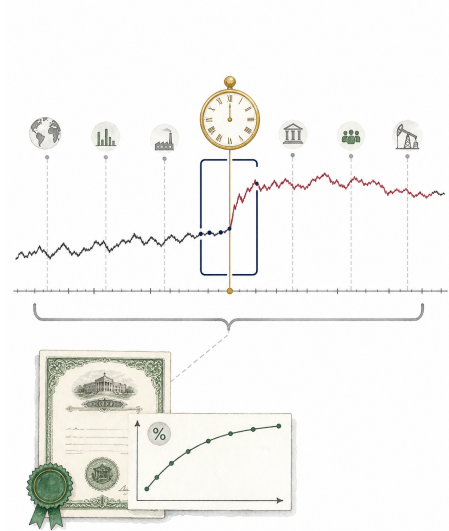

The event window compares prices immediately before and after the timestamp

$$\Delta A_i = A_{i,t^+} - A_{i,t^-}, \qquad TS_i = \text{futures-based surprise over } [t^-, t^+].$$

- Both variables use the same announcement event i and window.
- The short interval limits changes in slow-moving fundamentals.
- Identification comes from timing, not from adding many control variables.

A may be a yield, futures rate, equity price, or exchange rate. The transformation into a change or return must match the asset.

Narrow windows remove much of the competing news in a trading day

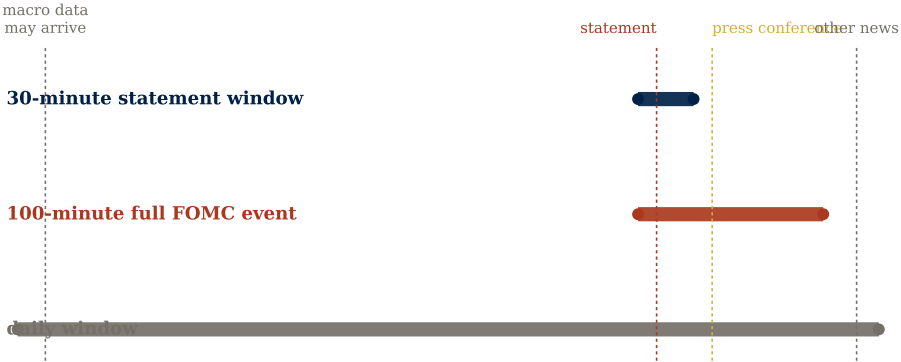

narrower windows remove more competing news but may retain market noise

The 30-minute statement window isolates the release; the full FOMC event includes the press conference; a daily window can contain macro data, earnings, geopolitical news, and market drift.

Window definitions follow FRBSF USMPD; daily contamination is schematic.

Yield changes and asset returns require different calculations

Asset	Event response	Natural unit
Interest rate or yield	$100(y_{t^+} - y_{t^-})$	basis points if y is in percentage points
Equity index	$100 \ln(P_{t^+}/P_{t^-})$	percent log return
Exchange rate	$100 \ln(S_{t^+}/S_{t^-})$	percent log return, with declared quote

> A coefficient is interpretable only when the response and surprise units are explicit.

P = asset price; S = exchange-rate quote. EURUSD and USDJPY have opposite base/quote conventions, so "currency appreciation" must be stated carefully.

The public database separates statement, press conference, full event, and minutes

USMPD sheet	Window	Information captured
Statements	−10 to +20 minutes	decision and statement wording
Press Conferences	−10 to +60 minutes	prepared remarks and questions
Monetary Events	statement −10 to press conference +60	full scheduled FOMC event
Minutes	−10 to +20 minutes	later release of meeting detail

Acosta et al. (2025), USMPD documentation; database page updated 3 Aug 2026.

USMPD = U.S. Monetary Policy Event-Study Database. Window endpoints are measured relative to each release or press-conference start time.

Appendix A14 · Data provenance

The basic regression asks how an asset responds per unit of surprise

$$\Delta A_i = \alpha + \beta TS_i + u_i.$$

- Each observation is one announcement event, not one trading minute.
- TS_i is the measured current-meeting surprise in bp.
- ΔA_i is the event-window asset response in its declared unit.
- u_i collects other event-specific price news and measurement error.

α = average response when the measured surprise is zero; β = response per 1 bp surprise; u_i = regression residual.

Appendix A10 · Inference Python P6 · Estimate regression

A slope of 0.41 means 10 bp of surprise predicts 4.1 bp in the yield

$$\widehat{\Delta y_i} = \widehat{\alpha} + 0.41\, TS_i, \qquad TS_i = 10 \text{ bp} \Longrightarrow \widehat{\Delta y_i} - \widehat{\alpha} = 4.1 \text{ bp}.$$

- The coefficient is not 41%; both axes are measured in bp.
- A positive slope means hawkish current-rate news tends to raise the yield.
- The response can be smaller than one because the two-year yield averages many expected future short rates.

The 0.41 value is the simple scheduled-statement slope in the frozen USMPD sample and is used for interpretation, not as a structural constant.

Across scheduled statements, positive target surprises raise two-year yields on average

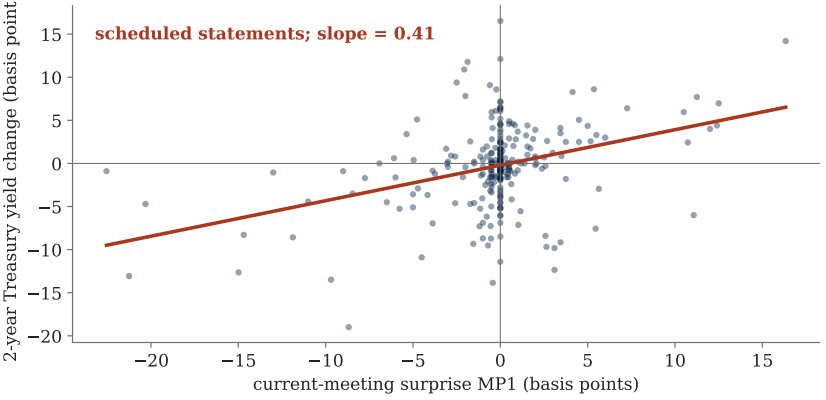

FRBSF USMPD, 1994–2026 scheduled statement windows; author's calculation.

The fitted line summarises conditional association across events. Dispersion around it reflects path news, information effects, market noise, and other event-specific components.

Python P6 · Reproduce regression

Anticipated actions should not explain the immediate market revision

Policy component	Known before release?	Expected immediate yield effect
Anticipated action	largely priced	approximately zero after conditioning on surprise
Unanticipated action	revealed at release	systematic change in short and intermediate yields

> Efficient-market logic predicts price response to new information, not to a policy action already embedded in prices.

Kuttner (2001), "Monetary policy surprises and interest rates."

"Approximately zero" is an empirical prediction subject to measurement error, path news, and imperfect pre-event pricing.

June and July 2026 changed the expected path in opposite directions

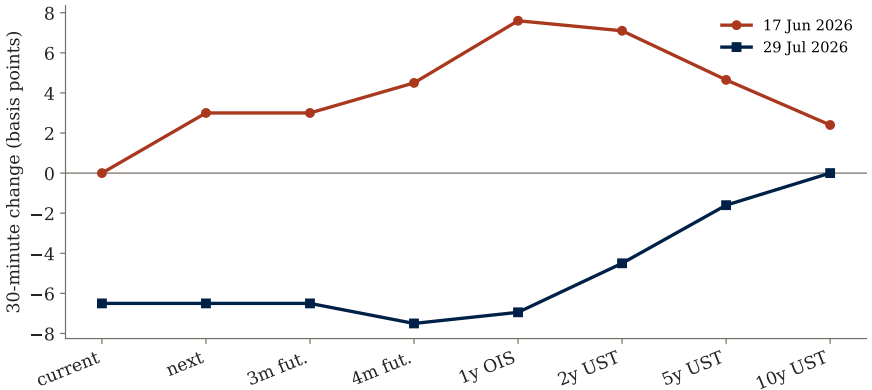

FRBSF USMPD, statement windows; futures, OIS, and Treasury changes converted to bp.

Both meetings maintained the target range. On 17 June, later-horizon rates rose; on 29 July, near-horizon rates fell. The curve response contains information beyond the current action.

Python P5 · Term-structure response

High-frequency identification requires fast prices and limited simultaneous news

- **Timing**: the release time is known precisely and observations bracket it.
- **Speed**: liquid prices incorporate public news quickly within the chosen window.
- **Exclusion**: no other major information arrives at the same timestamp.
- **Measurement**: the futures contract captures the intended surprise component.
- **Stability**: market functioning is adequate for observed prices to be informative.

The assumptions justify an event-time interpretation. They do not automatically identify a structural shock with all central-bank information removed.

A very narrow window can retain microstructure noise

- Bid–ask bounce can create a price change without a change in consensus value.
- Thin contracts may update slowly or trade only after the release.
- Different markets have different liquidity and timestamp precision.
- Press-conference answers arrive sequentially, so a longer window is economically necessary.

> Window choice balances contamination outside the event against measurement noise inside it.

Microstructure noise = short-lived price variation caused by trading mechanics rather than the underlying information value.

Timestamp association is strong evidence, but structural causality asks more

Claim	Event study supports	Additional requirement
Prices revised at release	strongly	precise timing and clean window
Rate surprise moved prices	conditionally	valid surprise measure
Pure policy shock caused response	not by timing alone	separate policy from central-bank information
Macro economy responds similarly	not directly	external transmission design and horizon

Structural shock = an interpretable disturbance isolated from central-bank information about economic fundamentals.

Appendix A12 · Information effect

Checkpoint: report window, surprise, response, and interpretation together

1. Name the event and exact release timestamp.
2. Declare the pre/post window for every market series.
3. Construct the surprise from an expectation-sensitive contract.
4. Transform each asset response into a stated unit.
5. Interpret the coefficient at a meaningful surprise size.
6. Separate timestamp evidence from stronger structural claims.

> A reproducible event study is a sequence of timed measurements, not a chart of prices "around policy."

Separate actions from words

What changes when the current target does not?

The statement and press conference can revise the expected path, risks, and economic outlook.

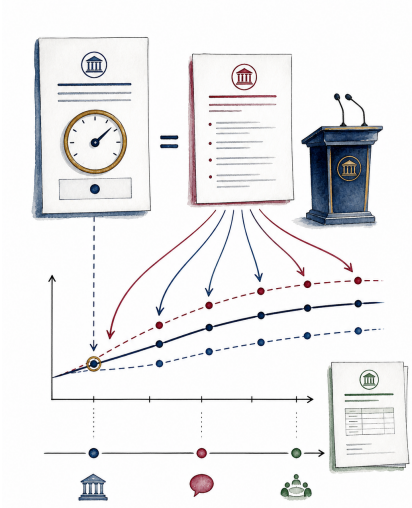

One FOMC event contains several communication channels

Channel	Timing	Information added
Policy statement	14:00 ET	current action, economic assessment, risks, vote
Press conference	from 14:30 ET	explanation, conditionality, questions, emphasis
Minutes	about three weeks later	discussion, alternatives, disagreements, detail
Projections / SEP	selected meetings	participant forecasts and policy-rate projections

Federal Reserve FOMC calendar and statement documentation.

SEP = Summary of Economic Projections. The channels have different timestamps, so a single daily return cannot assign news to each one.

The current decision and the expected path can move separately

$$\underbrace{\Delta \mathbf{f}_i}_{\text{vector of futures-rate changes}} = \underbrace{\mathbf{b}_T T_i}_{\text{current-target direction}} + \underbrace{\mathbf{b}_P P_i}_{\text{future-path direction}} + \underbrace{\varepsilon_i}_{\text{other variation}} .$$

- T_i is strongest in contracts closest to the current meeting.
- P_i loads more strongly on later contracts.
- A zero current action can coexist with a large P_i.

$\Delta \mathbf{f}_i$ = changes across futures horizons; $\mathbf{b}_T, \mathbf{b}_P$ = response patterns; T_i, P_i = target and path news.

Appendix A11 · Factor geometry

The yield-curve response indicates where expectations changed

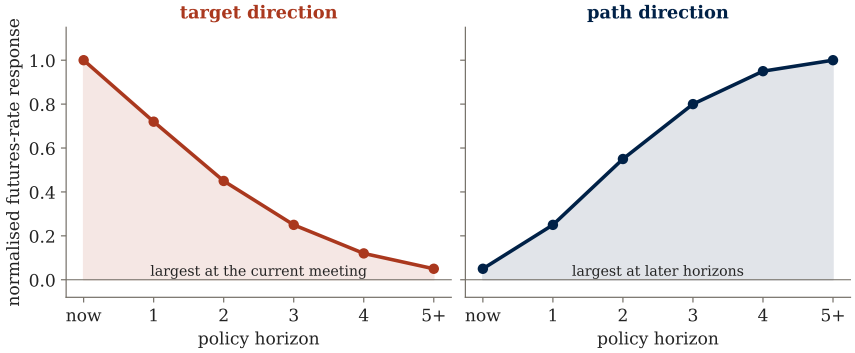

- A front-loaded response points toward current-meeting news.
- A response concentrated farther out points toward guidance or outlook news.

Normalised conceptual directions, not estimated June or July values. Factor loadings depend on the contracts and sample.

Appendix A11 · Normalisation

June and July maintained the target but sent opposite path signals

Statement date	MP1	MP2	2-year Treasury
17 Jun 2026	0.0 bp	+3.0 bp	+7.1 bp
29 Jul 2026	−6.5 bp	−6.5 bp	−4.5 bp

- Both statements maintained the 3.50–3.75% target range.
- June raised near-term and two-year rates; July lowered them.
- The current action alone would label both events identically.

Federal Reserve statements; FRBSF USMPD statement windows.

MP2 measures the expected rate at the next scheduled meeting; the two-year yield incorporates a much longer sequence of expected rates and term premia.

The target and path are two directions in a term-structure response

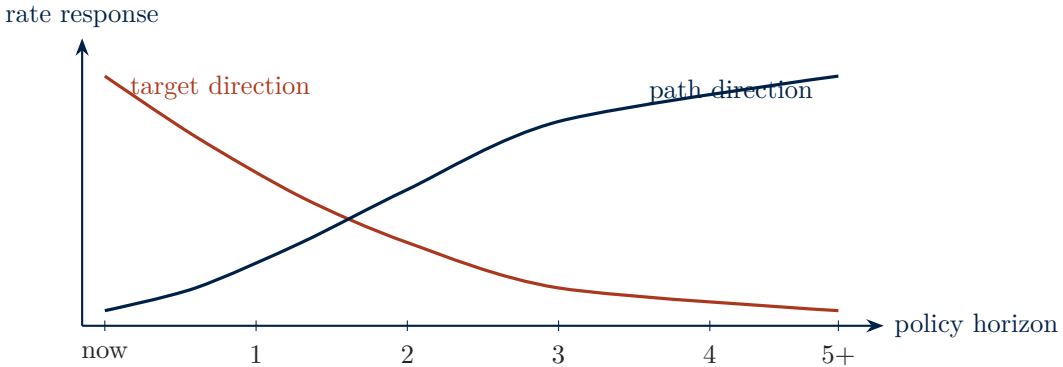

A factor is a common pattern across several rate changes. Its sign and scale must be normalised before "positive" can be called hawkish or dovish.

One surprise factor is not enough to explain FOMC yield-curve movements

- Current-meeting futures isolate a target component.
- Later futures often move when the target surprise is near zero.
- A second "path" component summarises common revisions at later horizons.
- Longer Treasury yields respond substantially to path news.

> The policy announcement is multidimensional because it changes both today's instrument and beliefs about future policy.

Gurkaynak, Sack, and Swanson (2005), "Do Actions Speak Louder Than Words?"

The main lecture uses target/path intuition. Principal-components estimation and rotation choices remain in the technical appendix.

Appendix A11 · Estimation details

Communication news is the rate revision not exhausted by the current target

$$\text{communication component at horizon } h \approx \Delta f_i^{(h)} - \widehat{b}_h \, TS_i.$$

- TS_i captures the current-meeting component.
- $\widehat{b}_h TS_i$ is the fitted change at horizon h associated with that component.
- The residual can contain guidance, outlook news, risk reassessment, and measurement error.

$\Delta f_i^{(h)}$ = futures-rate change at horizon h; $\widehat{b}_h$ = estimated target loading. "Communication" is an interpretation requiring validation.

Press conferences now contribute price discovery comparable to statements

USMPD window, 2019–2026	Events	Mean |Δ 2y UST|	Reading
Statement	67	3.22 bp	decision and written wording
Press conference	62	4.38 bp	explanation and questions
Full monetary event	67	5.43 bp	combined net revision

> The press conference is part of the monetary event, not postscript commentary.

FRBSF USMPD; author's calculation using absolute two-year Treasury yield changes.

Mean absolute change measures price-discovery magnitude, not effectiveness or causal importance. Samples differ slightly because some events lack press conferences.

Central-bank information can alter the usual rates–equities pattern

Joint response	Simple reading	Alternative information reading
Yields up, equities down	tighter expected policy raises discount rates	standard policy-news pattern
Yields up, equities up	puzzling under discount rates alone	bank revealed stronger growth or inflation outlook
Yields down, equities down	puzzling under easing alone	bank revealed weaker fundamentals

A central-bank information effect occurs when the announcement changes private beliefs about economic fundamentals as well as policy. Joint asset responses are diagnostic, not definitive.

Appendix A12 · Structural distinction

29 July shows why a word count cannot replace the market surprise

- The Committee maintained the target range by a 9–3 vote.
- Three dissenters preferred a 25 bp increase at that meeting.
- The statement said inflation remained elevated and emphasised price stability.
- Yet MP1 and MP2 both fell by 6.5 bp in the statement window.

> Hawkish phrases can coexist with a dovish market revision because prices respond to the whole statement relative to expectations.

Federal Reserve statement, 29 Jul 2026; FRBSF USMPD.

The vote and phrases are public text features. The surprise is the market's before/after rate revision; they measure different objects.

Checkpoint: ask which horizon and which channel moved

1. Separate current-target news from revisions to later policy rates.
2. Inspect a vector of futures, OIS, and Treasury changes across horizons.
3. Distinguish statement, press-conference, full-event, and minutes windows.
4. Use joint asset responses to diagnose central-bank information.
5. Treat wording as a candidate explanation that requires economic validation.

> Communication matters when it changes beliefs. Market prices show whether those beliefs changed.

PART 6 OF 6

Turn words into a measurable signal

Can text explain what markets learned?

Define the economic object, use an auditable baseline, and validate it against prices rather than intuition.

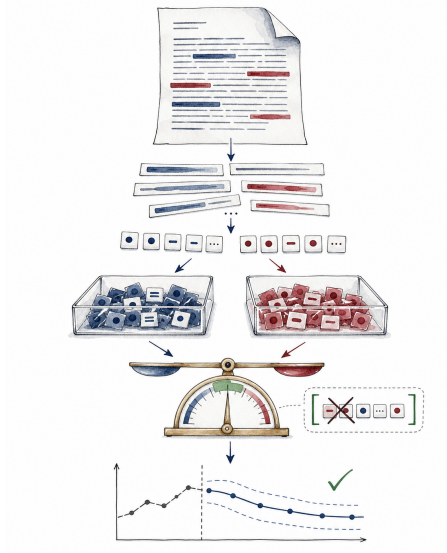

Text fills the missing middle between announcement and price revision

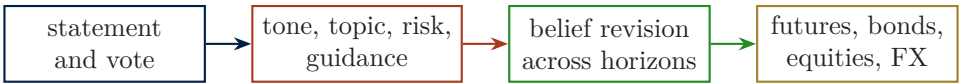

> Market prices measure whether beliefs changed; text features help describe which public content may have changed them.

A text feature is a reproducible numerical property of a document, such as a phrase count, topic share, uncertainty score, or semantic similarity.

Define the economic object before selecting the text method

Object	Question	Example signal
Policy tone	higher or lower expected rates?	hawkish/dovish phrases
Topic	what receives attention?	inflation, labour, financial stability
Uncertainty	how confident is the assessment?	uncertainty and risk language
Guidance	how is future policy conditioned?	dates, thresholds, data dependence
Disagreement	how unified is the committee?	dissents, alternatives, attribution

A higher topic share means more text devoted to a subject; it does not by itself mean the central bank expects that variable to rise.

Match the document to the question and release time

- **Statement**: concise decision, assessment, risks, and vote at the policy timestamp.
- **Press conference**: richer explanation and unscripted answers within the same monetary event.
- **Minutes**: detailed discussion released later; inappropriate for explaining the original statement-window move.
- **Speeches**: individual views, different audiences, and endogenous speaking dates.
- **Reports**: broad analysis with slower publication frequency and changing length.

Corpus = the declared set of documents analysed. Mixing statements, speeches, and minutes without document controls changes both language and institutional meaning.

A transparent dictionary is a reproducible baseline

Hawkish candidates	Dovish candidates
inflation remains elevated	downside risks
price stability	labour-market conditions have cooled
raise, strong, solid	lower, weakened, slowed

- Every term is visible and can be challenged.
- Phrase matching preserves more context than isolated words.
- Domain-specific validation is still required.

Dictionary or lexicon = a declared mapping from words or phrases to categories. Token = one extracted word unit used by the counting rule.

Python P7 · Dictionary score

Pre-processing choices determine what the dictionary counts

1. Extract only the official statement body; exclude menus, media contacts, and implementation links.
2. Standardise case, whitespace, apostrophes, and hyphen variants.
3. Match multi-word phrases before single words.
4. Remove matched phrases before counting their component words.
5. Retain document date, source URL, word count, and term-level contributions.

> Reproducibility begins with document extraction, not with the final sentiment formula.

Pre-processing = deterministic steps that convert raw document text into units used by the measurement rule.

Appendix A13 · Text audit

Normalising by length makes documents of different size comparable

$$\text{net tone}_i = 100 \times \frac{H_i - D_i}{W_i}.$$

- H_i is the number of hawkish dictionary matches.
- D_i is the number of dovish dictionary matches.
- W_i is total extracted word count.
- The result is net hawkish matches per 100 words.

The denominator controls document length, not linguistic importance. A short phrase and a long sentence can receive the same count.

Appendix A13 · Formula audit Python P7 · Calculate score

Negation and context can reverse the label attached to a word

Sentence fragment	Dictionary problem
"not appropriate to raise rates"	isolated "raise" is falsely hawkish
"inflation is less elevated"	"elevated" ignores direction of change
"risks to employment have diminished"	risk word alone misses improved assessment
"would raise if inflation reaccelerates"	conditional statement is not an action

> A dictionary is a baseline whose errors should be visible; adding complexity is justified only when it fixes documented failures.

Negation window = a rule that checks nearby words such as "not" or "less"; it cannot resolve all grammar or conditionality.

Official 2026 statements produce a small, auditable tone comparison

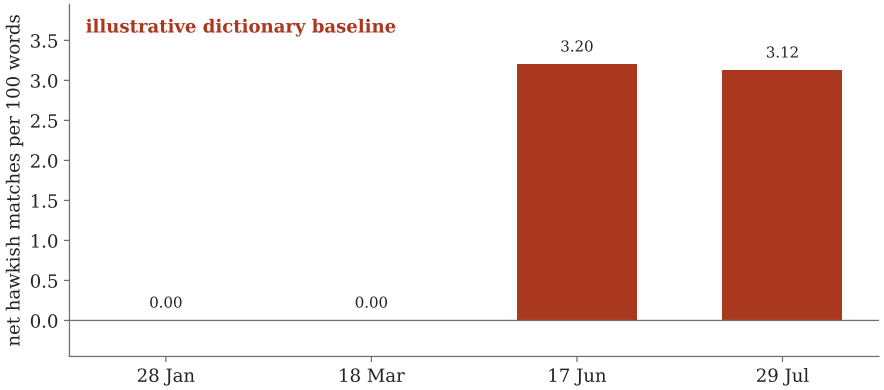

Official FOMC statements; illustrative dictionary declared on the previous slides.

The June and July statements are shorter and use more matched hawkish phrases. The score describes this dictionary, not an official Federal Reserve tone series.

Python P7 · Reproduce score

Term-level contributions explain every aggregate score

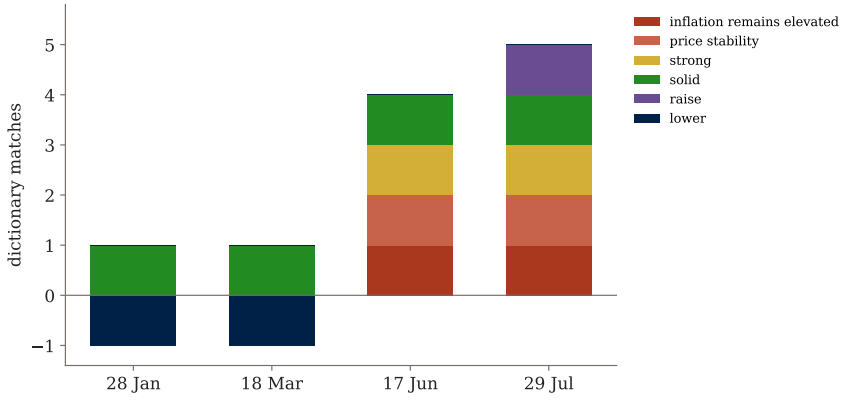

Official FOMC statements; author's transparent phrase-count calculation.

Positive bars are hawkish matches; negative bars are dovish matches. "Lower" in January and March comes from dissents preferring a lower target, illustrating why attribution matters.

Appendix A13 · Attribution and scope

Validate tone against a future-rate revision, not against intuition alone

Date	Net tone per 100 words	MP2 surprise	Same direction?
28 Jan	0.00	0.0 bp	uninformative
18 Mar	0.00	−0.5 bp	uninformative
17 Jun	+3.20	+3.0 bp	yes
29 Jul	+3.13	−6.5 bp	no

> A valid tone measure should explain a declared market object out of sample, not merely sound plausible when read after the event.

MP2 = expected-rate surprise at the next scheduled meeting. The July mismatch is evidence to inspect context and expectations, not permission to relabel the market response.

Four statements cannot establish a causal relation between tone and prices

- The sample is too small for reliable correlation or regression inference.
- Statement format and length changed during 2026.
- Dictionaries can be selected after seeing market reactions.
- Policy action, economic news, and wording arrive together.
- A credible study needs a longer frozen corpus, pre-declared rules, and out-of-sample validation.

> The four-event exercise is a computation audit and counterexample generator, not an empirical estimate of a language effect.

Out-of-sample = documents not used to choose the dictionary or model. Pre-declared = rules fixed before examining the target outcomes.

A text score is a measurement, not a policy preference or structural shock

It can measure	It cannot establish alone
matched language frequency	what the committee privately believes
document-level differences	whether communication caused the market move
reproducible tone under a declared lexicon	whether policy is objectively too tight or too loose
inputs for validation and forecasting	a pure structural monetary-policy shock

A hawkish score describes language associated with higher expected rates under the declared rule; it is not approval, criticism, or a welfare judgement.

The complete workflow links the financial contract to the document

- The contract supplies the expectation baseline and surprise.
- The timestamp supplies high-frequency identification.
- The document supplies interpretable candidate signals.
- Validation determines which signals deserve economic interpretation.

Each arrow adds assumptions. Preserve the intermediate data so every transformation can be audited.

Final checklist: state exactly what was measured and what can be claimed

1. Define the expectation object, horizon, instrument, and pre-event timestamp.
2. Decode settlement and remove meeting-month calendar weights.
3. Calculate the before/after surprise with a sign and unit audit.
4. Match asset responses to the same event window.
5. Separate current-target, path, and economic-information components.
6. Extract transparent text features from the correct document.
7. Validate out of sample and keep structural claims narrower than the evidence.

> Central-bank news is not the announced rate. It is the market's measured revision to beliefs at the announcement timestamp.

References and source notes

- Kuttner, K. N. (2001), "Monetary Policy Surprises and Interest Rates," *Journal of Monetary Economics* 47, 523–544.
- Gurkaynak, R. S., Sack, B., and Swanson, E. (2005), "Do Actions Speak Louder Than Words?" *International Journal of Central Banking* 1, 55–93.
- Jarocinski, M., and Karadi, P. (2020), "Deconstructing Monetary Policy Surprises," *American Economic Journal: Macroeconomics* 12, 1–43.
- Bauer, M. D., and Swanson, E. T. (2023), "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," *NBER Macroeconomics Annual* 37.
- Acosta, M., et al. (2025), U.S. Monetary Policy Event-Study Database; Federal Reserve Bank of San Francisco.
- Federal Reserve, 2026 FOMC statements and calendar; Federal Reserve Bank of New York, EFFR methodology; CME, 30-Day Federal Funds futures contract specifications.
- Kansoy, F., Oxford Macro-Finance 2026, `ml1.html`, `ml2.html`, and monetary-policy-surprise modules: local sequencing, examples, and extensions adapted for this course.

All 2026 event values use the frozen USMPD file updated 3 Aug 2026. Text scores are disclosed classroom calculations from official statement pages.

Use, attribution, and course website

- Course materials for the Computational Finance / FinTech Summer School.
- You may use them for personal study and course work.
- Do not redistribute, publish, or adapt the complete deck without the author's permission.
- Cite the named papers, central banks, exchange, and data providers when using their material.
- Course materials and updates: www.fatihkansoy.com

> Contact the author about reuse or corrections before publication.

Technical appendix map

A1 Notation and units

A2 EFFR median

A3 Futures settlement

A4 Meeting-month rate

A5 Kuttner derivation

A6 Sign and bp audit

A7 Late-month case

A8 Q versus P

A9 Event returns

A10 Regression inference

A11 Target and path

A12 Information effect

A13 Text audit

A14 Data provenance

> The appendix holds derivations, qualifications, and audit trails that support the main measurement sequence.

Python practical map

A1. Notation and unit conventions

Symbol	Meaning	Unit
F_t	30-Day Federal Funds futures price	price points
$\widetilde{f}_t$	$100 - F_t$, implied monthly-average rate	annual percentage points
N, d	days in month; announcement day	calendar days
TS_i	current-meeting target surprise	percentage points or bp
Δy_i	event-window yield change	percentage points or bp
R_i	event-window asset return	percent

Conversion: 1 percentage point = 100 bp, 1 bp = 0.01 percentage point.

← **Return to three objects** Appendix map

A2. The EFFR is a volume-weighted median, not a simple average

Transaction rate	Volume	Cumulative volume	Median test
5.30%	20	20	below 50
5.31%	40	60	crosses 50
5.32%	25	85	above 50
5.33%	15	100	above 50

> The 50th percentile of transaction volume is reached at 5.31%, so the EFFR is 5.31% in this illustrative distribution.

Volume-weighted median = rate at which cumulative eligible transaction volume reaches 50% of total volume. The example volumes are illustrative.

Federal Reserve Bank of New York, Reference Rates methodology.

← **Return to three rates**

Appendix map

A3. Futures settlement maps realised daily rates into one final price

$$\bar{r}_m = \frac{1}{N_m} \sum_{j=1}^{N_m} r_{m,j},$$

$$F_m^{\text{settle}} = 100 - \bar{r}_m,$$

$$\text{contract cash value} = \text{exchange multiplier} \times F_m^{\text{settle}}.$$

- The daily EFFR is repeated for calendar days governed by the contract convention.
- The contract is cash settled; no overnight loan is delivered.
- Intraday prices before settlement reflect the market's valuation of the unknown monthly average.

Cash settlement = final gains and losses are paid in money using the settlement index rather than delivery of an underlying loan.

← **Return to contract object**

Appendix map

A4. Derive the expected post-meeting rate from the monthly average

$$\mathbb{E}_t[\bar{r}_m] = \frac{1}{N}\left[(d-1)r_{\text{pre}} + (N-d+1)\mathbb{E}_t[r_{\text{post}}]\right],$$

$$N\mathbb{E}_t[\bar{r}_m] - (d-1)r_{\text{pre}} = (N-d+1)\mathbb{E}_t[r_{\text{post}}],$$

$$\mathbb{E}_t[r_{\text{post}}] = \frac{N\mathbb{E}_t[\bar{r}_m] - (d-1)r_{\text{pre}}}{N-d+1}.$$

The derivation assumes one average pre-meeting rate and one average post-meeting rate. Multiple implementation dates require the calendar sum to be split further.

← **Return to post-meeting rate**

Appendix map

A5. The Kuttner scaling factor follows from differencing the identity

Hold the realised pre-meeting contribution fixed across a short event window:

$$\Delta\mathbb{E}[\bar{r}_m] = \frac{N-d+1}{N}\,\Delta\mathbb{E}[r_{\text{post}}],$$
$$\Delta\mathbb{E}[r_{\text{post}}] = \frac{N}{N-d+1}\,\Delta\mathbb{E}[\bar{r}_m],$$
$$TS = \frac{N}{N-d+1}\,\Delta(100-F) = -\frac{N}{N-d+1}\,\Delta F.$$

The event-window assumption is that only the expected post-meeting component changes at the announcement; pre-meeting realised rates do not.

← **Return to surprise formula**

Appendix map

A6. Sign and basis-point audit for the worked surprise

Step	Calculation	Result
Futures-price change	$94.725 - 94.775$	-0.050
Implied monthly-rate change	$-(-0.050)$	$+0.050$ pp
Calendar scaling	$(30/15)(+0.050)$	$+0.100$ pp
Basis-point conversion	$100(+0.100)$	$+10$ bp

> Negative futures-price move → positive implied-rate move → hawkish rate surprise.

pp = percentage points. Do not multiply by 10,000 when the input rate is already expressed in percentage points.

← **Return to worked surprise**

Appendix map

A7. Late-month meetings shift information into the following contract

- On day 30 of a 30-day month, the current contract gives the new rate only 1/30 weight.
- Its scaling factor is 30, so a 0.5 bp quote error becomes a 15 bp surprise error.
- The next-month contract places nearly full weight on the post-meeting rate.
- A combined estimator uses both contracts and the exact expected rate path across the boundary.

$$\widehat{TS} = w\,\widehat{TS}_{\text{current}} + (1-w)\,\widehat{TS}_{\text{next}},$$

w is chosen from the calendar and measurement precision; there is no universal fixed weight for every meeting.

← **Return to late-month problem**

Appendix map

A8. Pricing probabilities Q need not equal physical probabilities P

$$\text{futures-implied rate} = \mathbb{E}_t^Q[r_T], \qquad \text{forecast belief} = \mathbb{E}_t^P[r_T].$$

- Under P, states are weighted by their expected real-world frequency.
- Under Q, states are also weighted by how marginal wealth is valued in each state.
- At very short horizons, the wedge may be small but should not be assumed away silently.
- Survey forecasts can help assess whether a persistent pricing wedge is material.

P and Q are probability measures, not competing claims that one set of prices is "wrong."

← **Return to probability caveat** Appendix map

A9. Event returns are additive approximations only over short moves

$$R_i^{\log} = 100 \ln\left(\frac{P_{i,t^+}}{P_{i,t^-}}\right),$$

$$R_i^{\text{simple}} = 100\left(\frac{P_{i,t^+}}{P_{i,t^-}} - 1\right),$$

$$R_i^{\log} \approx R_i^{\text{simple}} \quad \text{for small price changes.}$$

- Yield changes are differences, not returns on a bond portfolio.
- A price return and a yield change naturally have opposite signs for a fixed-cash-flow bond.

A bond yield increase generally implies a bond-price decline, but duration is needed to convert yield changes into approximate bond returns.

← **Return to response units**

Appendix map

A10. Event-study inference should allow unequal residual variance

$$\widehat{\beta} = \frac{\sum_i (TS_i - \overline{TS})(\Delta A_i - \overline{\Delta A})}{\sum_i (TS_i - \overline{TS})^2}.$$

- Large-surprise days may have larger residual variance than quiet days.
- Heteroskedasticity-robust standard errors protect inference against unequal variance.
- A few crisis events can dominate the slope; influence and subsample checks matter.
- Measurement error in TS_i can attenuate the estimated coefficient.

Heteroskedasticity = residual variance differs across events. Attenuation = classical regressor measurement error biases the slope toward zero.

← **Return to regression** Appendix map

A11. Target and path factors require a sign and scale normalisation

$$\Delta \mathbf{f}_i = \mathbf{B}\mathbf{s}_i + \boldsymbol{\varepsilon}_i, \qquad \mathbf{s}_i = (T_i, P_i)'.$$

- Principal components find directions that explain common variance across contract changes.
- Multiplying a factor and its loadings by -1 leaves fitted values unchanged.
- The target factor is normalised to align with the current-meeting surprise.
- The path factor is rotated or constrained to have little effect on the current target and stronger later loadings.

$\mathbf{B}$ = matrix of horizon loadings; $\mathbf{s}_i$ = event factor scores. Economic names come from normalisation and response patterns, not from PCA alone.

← **Return to target and path** Appendix map

A12. A central-bank information effect mixes policy and outlook news

$$\Delta y_i = a_y \,\text{policy}_i + b_y \,\text{information}_i + u_{y,i},$$
$$R_i^{\text{equity}} = a_s \,\text{policy}_i + b_s \,\text{information}_i + u_{s,i}.$$

- A hawkish policy component tends to raise yields and lower equities through discount rates.
- Positive growth information can raise both yields and equities.
- Sign restrictions use these joint responses to separate the two components.

Jarocinski and Karadi (2020); Bauer and Swanson (2023).

Sign restrictions are identifying assumptions. They improve interpretation but do not make the components directly observed.

← **Return to information effect**

Appendix map

A13. Text-score audit: extraction, attribution, denominator, and context

1. Freeze the official source HTML and store its URL and date.
2. Extract the statement body while retaining voting paragraphs separately if attribution matters.
3. Count phrases before words and publish the complete lexicon.
4. Report H_i, D_i, W_i, the normalised score, and every term contribution.
5. Inspect negation, conditional language, and speaker attribution by hand on a validation sample.
6. Freeze rules before testing them against market surprises.

The Day 7 baseline intentionally counts voting language; its "lower" matches in January and March demonstrate the resulting attribution problem.

← **Return to contributions**

Appendix map

A14. Frozen data provenance, windows, and units

File / series	Source	Convention
USMPD workbook	FRBSF, updated 3 Aug 2026	statement $-10/+20$ min; press conference $-10/+60$ min
MP1 / MP2	USMPD	percentage-point rate changes; multiply by 100 for bp
Futures, OIS, UST	USMPD / LSEG Tick History	rate or yield changes in percentage points
S&P 500, DXY, FX	USMPD / LSEG Tick History	percent returns
2026 statement HTML	Federal Reserve	official release body and voting paragraph
Text score	disclosed Day 7 calculation	net dictionary matches per 100 words

Derived CSV files preserve source columns and dates. The notebook does not download live data, so results remain reproducible after the course.

← **Return to USMPD windows**

Appendix map

Python practical map

P1 Load and inspect data

P2 Futures price to rate

P3 Post-meeting rate

P4 Target surprise

P5 Recent event charts

P6 Event regression

P7 Statement text score

P8 Validation checks

- The notebook reproduces calculations already developed in the lecture.
- Data are frozen locally with source, window, and unit metadata.
- Each stage includes an expected output or numerical assertion.

Technical appendix map

P1. Load the frozen statement-window data

```python
import pandas as pd

data_path = "../data/derived/usmpd_statement_windows.csv"
events = pd.read_csv(data_path, parse_dates=["date"])

# Inspect rows, columns, and the latest event before calculating
print(events.shape)
print(events[["date", "MP1", "MP2", "UST2Y"]].tail())

```

Expected shape: 278 events and 27 columns in the frozen subset. Interest-rate changes are in percentage points.

← Return to data windows

Practical map

P2. Convert a futures price into an implied monthly rate

```python
def implied_monthly_rate(futures_price):
    """Return the annual percentage-point rate implied by the quote."""
    return 100.0 - futures_price

prices = [94.50, 95.00, 95.50]
rates = [implied_monthly_rate(price) for price in prices]

print(rates)
assert rates == [5.50, 5.00, 4.50]

```

The decreasing output checks the sign: a higher futures price implies a lower monthly-average EFFR.

← **Return to price–rate sign**

Practical map

P3. Recover the expected post-meeting rate

```python
def post_meeting_rate(monthly_rate, pre_rate, days, meeting_day):
    """Remove pre-meeting calendar days from the monthly average."""
    pre_days = meeting_day - 1
    post_days = days - meeting_day + 1
    return (days * monthly_rate - pre_days * pre_rate) / post_days

expected_post = post_meeting_rate(5.225, 5.10, 30, 16)
print(round(expected_post, 3))
assert abs(expected_post - 5.35) < 1e-12

```

Expected output: 5.350. The assertion verifies that the calendar-weight calculation reproduces the worked example.

← Return to worked rate

Practical map

P4. Calculate a target surprise and audit its sign

```python
def target_surprise_bp(price_before, price_after, days, meeting_day):
    """Return the calendar-scaled current-meeting surprise in bp."""
    post_days = days - meeting_day + 1
    price_change = price_after - price_before
    surprise_pp = -(days / post_days) * price_change
    return 100.0 * surprise_pp

surprise = target_surprise_bp(94.775, 94.725, 30, 16)
print(round(surprise, 1))
assert abs(surprise - 10.0) < 1e-10

```

Expected output: +10.0 bp. The leading minus sign converts the futures-price fall into a positive rate surprise.

← Return to worked surprise

Practical map

P5. Compare recent surprise and yield-curve responses

```python
recent = events.loc[events["date"].dt.year == 2026].copy()

# USMPD rate changes are percentage points; convert them to bp once
rate_columns = ["MP1", "MP2", "UST2Y", "UST10Y"]
recent[[f"{c}_bp" for c in rate_columns]] = 100 * recent[rate_columns]

print(recent[["date", "MP1_bp", "MP2_bp", "UST2Y_bp"]])

```

The latest row should show -6.5 bp for MP1, -6.5 bp for MP2, and -4.5 bp for the two-year Treasury yield.

← Return to recent surprises

Practical map

P6. Estimate the event-study slope with transparent arrays

```python
import numpy as np

sample = events.loc[
    (events["Unscheduled"] == 0) & events["MP1"].notna()
    & events["UST2Y"].notna()
].copy()

x = 100 * sample["MP1"].to_numpy()        # surprise in bp
y = 100 * sample["UST2Y"].to_numpy()      # yield change in bp
slope, intercept = np.polyfit(x, y, 1)

print("slope:", round(slope, 2))

```

Expected slope: about 0.41 bp of two-year yield change per 1 bp current-meeting surprise. Robust inference is an extension beyond this calculation.

← **Return to event scatter**

Practical map

P7. Calculate an auditable statement tone score

```python
import pandas as pd

statements = pd.read_csv("../data/derived/fomc_2026_statements.csv")

# Counts were produced by the disclosed phrase-first dictionary
statements["tone_check"] = 100 * (
    statements["hawkish_count"] - statements["dovish_count"]
) / statements["word_count"]

print(statements[["date", "word_count", "tone_check"]])
assert statements["tone_check"].equals(statements["net_tone_per_100_words"])

```

The equality assertion verifies the aggregation formula. The contribution file shows which phrases produced each count.

← Return to tone scores

Practical map

P8. Use assertions to protect units and accounting identities

```python
# 1. Futures price and implied rate must sum to 100
assert abs(94.775 + implied_monthly_rate(94.775) - 100) < 1e-12

# 2. Rebuild the monthly average from the recovered post rate
rebuilt = (15 * 5.10 + 15 * expected_post) / 30
assert abs(rebuilt - 5.225) < 1e-12

# 3. Percentage points to basis points is multiplication by 100
assert 100 * (-0.065) == -6.5

print("all checks passed")

```

Assertions stop the notebook when a financial identity or unit conversion fails; they are part of the analysis, not optional decoration.

← **Return to full workflow**

Practical map