Measuring Expectations

Surveys, markets, projections, text, and the wedges behind each measure

Macro-Finance · Lecture 2 · Trinity Term 2026

Fatih Kansoy
University of Oxford · Saïd Business School · Department of Economics

Three empirical steps turn beliefs into evidence

02A
Define the latent belief, instrument, and wedge

→

02B
Read expectations from futures and asset prices

→

02C
Use event timing for causal identification

$$
\begin{aligned}
m_{k,t} &= g_k(\theta_t) + w_{k,t} + \varepsilon_{k,t} \\
f_{t,m} &= 100 - F_{t,m} \\
\Delta_{\mathcal{W}} x_e & \quad \text{isolates event-time news.}
\end{aligned}
$$

A measure becomes useful only after we identify its economic object and its wedge.

Part 02A

Fundamentals

Latent beliefs, observed instruments, probability measures, and validation

Every core macro equation asks what agents expect next

Household choice
Expected consumption value and real returns govern saving.

Borrowing and investment
Expected inflation converts nominal rates into real rates.

Price-setting
Expected future inflation enters today's pricing decision.

Long rates
Expected short rates create the policy path embedded in yields.

> **Expectations are structural objects; measurement determines whether the model can meet the data.**

The Euler equation links current spending to expected marginal utility

$$u'(c_t) = \beta(1 + r_t)\mathbb{E}_t\left[u'(c_{t+1})\right].$$

c_t: consumption · r_t: real return · β: discount factor.

Households postpone spending when the expected value of tomorrow's consumption, adjusted by the real return, is high enough relative to today.

Current saving depends on an expectation we cannot observe directly.

The Fisher relation makes expected inflation a real-rate input

$$r_t \approx i_t - \mathbb{E}_t[\pi_{t+1}].$$

r_t: real rate · i_t: nominal rate · $\mathbb{E}_t[\pi_{t+1}]$: expected inflation.

Investment and borrowing respond to the real rate. The central bank observes the nominal rate, but must estimate the inflation expectation that converts it into the relevant price.

A bank that cannot measure expected inflation cannot know the real stance of policy.

Forward-looking price-setters transmit beliefs into inflation

$$\pi_t = \kappa \, \widetilde{mc}_t + \beta \mathbb{E}_t[\pi_{t+1}].$$

$\widetilde{mc}_t$: real marginal-cost gap · κ: slope of the New Keynesian Phillips curve.

Anchored expectations limit the persistence of inflation after a cost or supply disturbance because future inflation does not ratchet upward one-for-one with the current outcome.

Forward guidance works by moving the expected short-rate path

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

$i_t^{(n)}$: n-period yield · $\tau_t^{(n)}$: term premium.

Guidance can alter $\mathbb{E}_t[i_{t+j}]$ for $j \geq 1$ even when the overnight rate does not move today.

The language changes policy when it changes the expected path embedded in long rates.

The belief is hidden; instruments leave different traces

$$\theta_{a,t}^{(h)}(X).$$

Agent a's subjective distribution for X_{t+h}, formed at t.

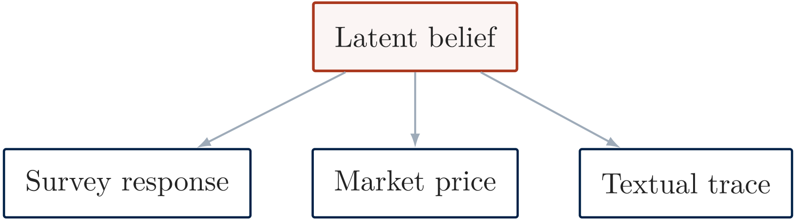

One latent belief leaves several observable traces

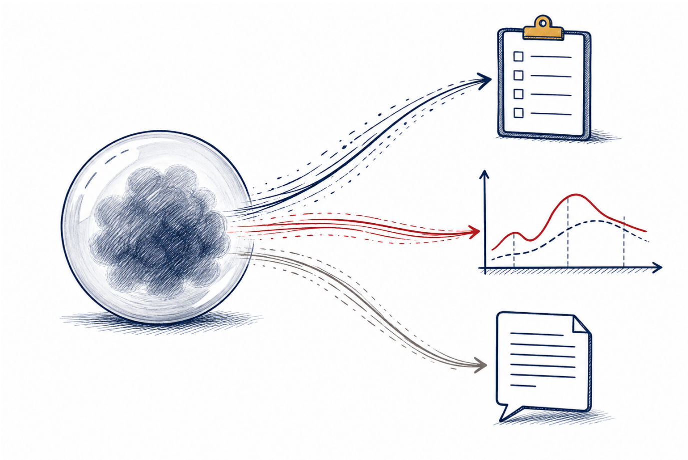

Editorial marker illustration generated for this lecture.

Every observed measure combines signal, wedge, and noise

$$m_{k,t}^{(h)} = g_k\left(\theta_{a,t}^{(h)}(X)\right) + w_{k,t}^{(h)} + \varepsilon_{k,t}^{(h)}.$$

- $g_k(\theta)$: moment or transformation captured by instrument k.
- $w_{k,t}^{(h)}$: structured, economically meaningful wedge.
- $\varepsilon_{k,t}^{(h)}$: sampling or measurement error.

The wedge is not classical noise; it is created by the instrument itself.

Three trace families create three characteristic wedges

Surveys
Direct, but sensitive to wording, framing, rounding, priming, and numeracy.

Markets
Continuous and traded, but contain premia, liquidity, collateral, and risk-neutral pricing.

Text
Timely salience and tone, but selection, strategic language, extraction error, and language drift.

Lucas (1976) sharpens the problem: when policy regimes change, agents revise beliefs, so measurement is part of identification.

US household surveys trade long histories for different objects

Michigan Survey of Consumers

Monthly since 1978. Point forecasts for the next 12 months and a five-year horizon. Clustering at 0, 5, and 10% can signal uncertainty and rounding.

NY Fed Survey of Consumer Expectations

Monthly panel since 2013. Probability bins recover individual means and uncertainty; repeated respondents reveal belief revision.

Point forecasts are easy to collect; distributions reveal uncertainty and heterogeneity.

UK and euro-area surveys add qualitative and experimental designs

BoE/Ipsos Inflation Attitudes Survey

Quarterly since 1999. Asks about "prices in the shops generally" and converts qualitative balances into a quantitative measure.

ECB Consumer Expectations Survey

Monthly panel since 2020. Distributional HICP elicitation and embedded randomised information treatments support causal research.

> **A survey can be a measurement instrument and an experimental platform.**

Survey design determines the expectation being measured

1. **Wording**: aggregate prices, inflation, and own living costs elicit different objects.
2. **Format**: a point, qualitative balance, and probability distribution are not interchangeable.
3. **Horizon**: cumulative, average annual, and point-in-time questions imply different timing.

> Apparent disagreement across surveys can come from question design rather than genuine belief differences.

Professional surveys supply long histories and distributions

- **Philadelphia Fed SPF**: US since 1968; probability bins recover uncertainty and disagreement. ECB SPF mirrors it from 1999.
- **Blue Chip Economic Indicators**: monthly forecasts from roughly 50 business economists since 1976.
- **Livingston Survey**: biannual forecasts since 1946 across banks, industry, and government.
- **BoE SEF / Consensus Economics**: UK external forecasts before MPC meetings and a broad institutional cross-section.

Professionals gain information but may herd

Incentives
Forecast accuracy is part of the professional's job.

Information
Real-time data, models, and market intelligence are available.

Herding
Reputation incentives can pull forecasts toward consensus.

Ang, Bekaert and Wei (2007), as reported in the source, find Michigan and SPF outperform the US term structure at one-year inflation forecasting.

Disagreement across populations is systematic, not noise

> Mankiw, Reis and Wolfers (2003): households, firms, and professionals can differ on inflation by percentage points, with larger gaps in high-inflation episodes.

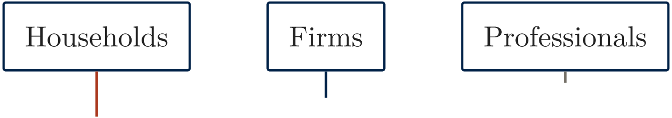

> **Before averaging beliefs, identify whose information, incentives, and attention generated them.**

A TIPS breakeven mixes expected inflation with two wedges

$$\mathrm{BE}_t^{(n)} = i_t^{(n)} - i_{t,\mathrm{TIPS}}^{(n)} = \mathbb{E}_t^Q[\pi_{t,t+n}] + \lambda_t + \ell_t.$$

- $\mathbb{E}_t^Q[\pi]$: risk-neutral expected inflation.
- λ_t: compensation for inflation risk.
- ℓ_t: liquidity and market-friction wedge.

A breakeven is compensation priced in bonds, not a pure physical expectation.

Inflation swaps remove the bond-liquidity wedge, not risk pricing

$$\mathrm{ILS}_t^{(n)} \approx \mathbb{E}_t^{Q}[\pi_{t,t+n}] + \lambda_t.$$

A fixed-rate payer receives realised CPI or HICP growth. Swaps avoid the TIPS-specific liquidity premium and are central to far-ahead anchoring analysis, including the euro-area 5y5y HICPxT rate.

Thin swap markets can still be illiquid, and Q-measure compensation remains distinct from average household or firm beliefs.

Jurisdiction determines the inflation index and market wedge

Market	Bond measure	Swap measure
United States	TIPS breakeven	CPI-U inflation swap
United Kingdom	ILG breakeven	RPI swap; CPI–RPI wedge
Euro area	OAT/Bund breakeven	HICPxT inflation swap

The label “market inflation expectation” hides the price index, contract design, liquidity, collateral, and risk compensation.

Fed funds futures price a monthly average of the overnight rate

$$F_{t,m} = 100 - \mathbb{E}_t^Q \left[\overline{r}_m^{EFFR} \right].$$

The contract settles against the arithmetic average EFFR in month m. It is therefore a clean instrument for near-term Fed expectations, but the meeting date determines how much of the month is exposed to the post-meeting rate.

The object is a monthly average, not the policy rate on one day.

OIS curves extend the expected policy path across tenors

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

The fixed leg exchanges against compounded SOFR, SONIA, or €STR. The curve reads out a sequence of expected overnight rates, plus a term-premium component.

Expectations about policy are a path, not a single number.

Options recover the distribution around the rate path

0–3 months

Fed funds futures: current-month averages, meeting probabilities, and target surprises.

Beyond three months

SOFR futures and OIS: a longer expected dollar or local-rate path.

Distribution tails

Caps, floors, and swaptions recover probabilities such as

$$\overset{Q}{\Pr}(i_{t+n} > \bar{i}),$$

not only the mean.

Population and jurisdiction define the empirical object

Population	United States	United Kingdom	Euro area
Households	Michigan; NY Fed SCE	BoE/Ipsos	ECB-CES
Firms	Atlanta BIE; Cleveland SoFIE	BoE DMP	Bundesbank; Bank of Italy; Ifo
Professionals	SPF; Blue Chip; Livingston	BoE SEF; Consensus	ECB SPF
Markets	FFF; TIPS; CPI swaps	SONIA; ILGs; RPI swaps	€STR; HICPxT; OAT/Bund

> **There is no generic "expected inflation" series: sampling frame and contract define the measure.**

Michigan, ILG breakevens, and ECB-CES are not substitutes

Michigan

US household point forecasts with rounding and attention wedges.

ILG breakeven

UK traded Q-measure with RPI and market-premium wedges.

ECB-CES

Euro-area household panel with HICP framing and distributions.

> A valid comparison holds fixed the agent, variable, horizon, moment, jurisdiction, and probability measure.

Survey responses contain design and reporting wedges

$$s_{i,t} = \theta_{i,t} + b(q_k, z_i) \\ + r_{i,t} + \varepsilon_{i,t}.$$

$b(q_k, z_i)$: question/framing effect · $r_{i,t}$: rounding/reporting component.

Bin design can prime answers near target; 0, 5, and 10% responses often reveal uncertainty rather than simple error.

Market prices convert physical beliefs into risk-neutral prices

$$\mathbb{E}_t^Q[X] = \mathbb{E}_t^P[X] + \rho_t^X,$$
$$\rho_t^X = \lambda_t + \ell_t + \chi_t.$$

λ_t: risk premium · ℓ_t: liquidity · χ_t: other market segmentation/collateral wedges.

The P–Q gap is not merely contamination: its variation is itself a macro-finance object.

Text reflects beliefs, strategy, extraction, and selection

$$
\begin{aligned}
T_t = {} & f(\theta_t, \text{strategy}_t) \\
& + \nu_t^{extract} + \nu_t^{select}.
\end{aligned}
$$

- Extraction depends on language, corpus, dictionary, model, and drift.
- Text producers are selected, not a random population.
- The advantage is real-time attention and narrative information.

Triangulate surveys, markets, and text to separate belief movement from instrument movement.

Triangulation asks whether different traces share one signal

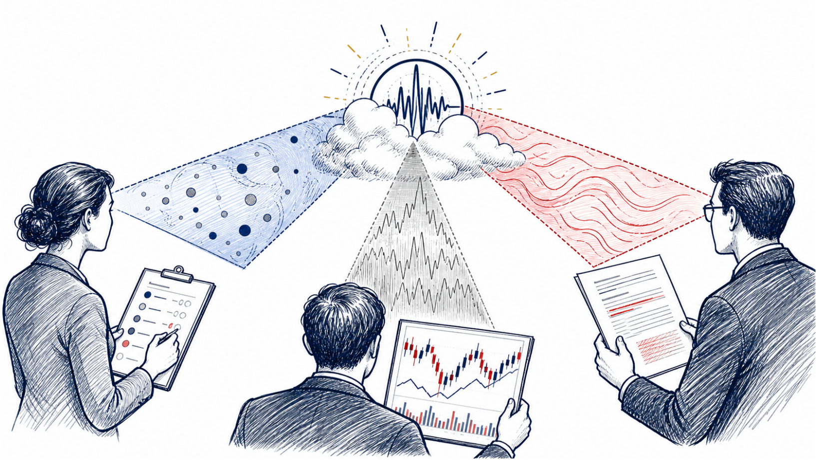

Editorial marker illustration generated for this lecture.

A five-year spot measure averages years one through five

$$\pi_t^{0,5} = \mathbb{E}_t\left[\frac{1}{5}\sum_{h=1}^{5}\pi_{t+h}\right].$$

Near-term energy prices, food shocks, taxes, and supply disruptions all enter this average. A rise therefore need not imply that the distant inflation anchor has moved.

A 5y5y forward measure isolates years six through ten

$$\pi_t^{5,5} = \mathbb{E}_t\left[\frac{1}{5}\sum_{h=6}^{10}\pi_{t+h}\right].$$

The forward window removes the first five years and asks whether inflation is expected to return near target after the current shock fades.

> Anchoring is a far-ahead statement; a high one-year or five-year spot measure is not proof that the anchor has moved.

The 2021–22 episode separated near-term inflation from the anchor

- Near-term measures rose sharply as current CPI, energy, and supply disruptions entered the first few forecast years.
- 5y5y measures moved less because they refer to inflation after the shock horizon.
- Central-bank benchmarks differ: US TIPS/CPI swaps, UK RPI swaps adjusted for the RPI–CPI wedge, and euro-area HICPxT swaps.

> **Always identify which years a horizon label actually averages.**

The same mean can hide very different inflation tails

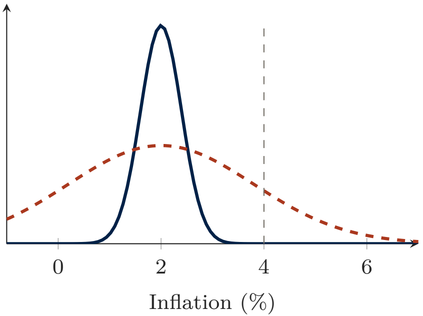

Agent A

$\mathbb{E}[\pi] = 2\%$, $\sigma = 0.4$ pp: mass is concentrated near target.

Agent B

$\mathbb{E}[\pi] = 2\%$, $\sigma = 1.8$ pp: $\Pr(\pi > 4\%) \approx 13\%$.

> **Anchoring is a property of the full distribution, not its centre alone.**

Distributional measures reveal risks before point forecasts move

Euro area, 2014–15

Option-implied distributions highlighted deflation tail risk that a mean forecast could understate.

US and UK, 2021–22

The option-implied probability of inflation above 4% rose before survey means fully adjusted.

Point forecasts are incomplete. Distributional surveys and option-implied densities are needed to measure anchoring and credibility risk.

The dot plot summarises individual policy assessments

$$\tilde{i}_{t+h}^{CB} = \mathrm{median}_j(i_{j,t+h}^{*,CB}).$$

- SEP dots report individual FOMC participants' appropriate-rate paths.
- The median is a focal summary; dispersion reveals internal disagreement.
- Dots are individual reaction functions, not a committee forecast.

The ECB SPF aggregates expert probability distributions

$$P_{t,h}(\pi \in B_b) = \frac{1}{N_t} \sum_i p_{i,b,t,h}.$$

The quarterly survey collects point and distributional forecasts for HICP inflation, GDP growth, and unemployment at one-, two-, and five-year horizons.

The SPF is central-bank data infrastructure, not a measure of the central bank's own belief.

Staff projections are conditional forecasts, not unconditional promises

$$\widehat{X}_{t+h}^{staff} = E_t^{staff}[X_{t+h} \mid \text{paths}].$$

The paths can include policy rates, exchange rates, fiscal policy, energy prices, and model assumptions. A gap from market expectations may therefore reflect different conditioning assumptions rather than different structural beliefs.

A central-bank projection can change the outcome it forecasts

$$\theta^{P}_{a,t+} = (1-\alpha_a)\theta^{P}_{a,t-} + \alpha_a \widehat{X}^{CB}_{t+h}.$$

α_a: weight agent a places on the official projection.

The central bank is both forecaster and actor. Publication can move private beliefs toward the policy-consistent path, changing the realised outcome used to evaluate the forecast.

Dot–market gaps can reveal disagreement or compensation

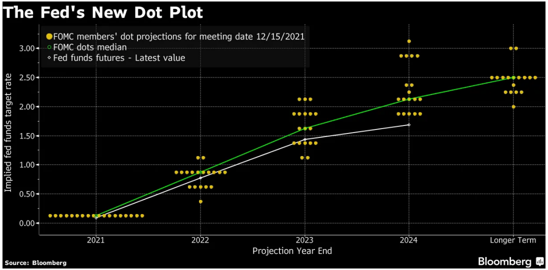

December 2021 SEP and fed funds futures.

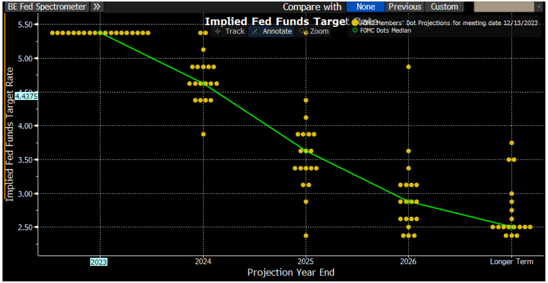

December 2023 SEP and implied target path.

A gap can reflect a different reaction function, risk compensation, or both.

Composite indices extract the signal shared across measures

$$m_{k,t} = \lambda_k f_t + u_{k,t}, \qquad f_t = \rho f_{t-1} + \eta_t.$$

The Fed Common Inflation Expectations index combines 21 survey and market indicators. Cross-series correlation identifies a latent component that no individual wedge-free measure can supply.

> High-variance series and structural breaks can distort factor loadings.

Text measures attention and narrative at high frequency

$$\frac{\#\{\text{inflation terms}\}}{\#\{\text{all terms}\}}.$$

- Central-bank text can separate supply- from demand-driven narratives.
- Earnings calls reveal firm-side concern about inflation and costs.
- Newspaper indices supply timely sentiment but depend on extraction.

Text often measures salience and framing before it measures a numerical belief.

Search intensity is public attention, not an expectation

$$Google_t(\text{“inflation”}).$$

Searches for inflation, price rises, and cost of living can lead survey releases and correlate with Michigan or SCE at weekly frequency.

> A search shows that inflation is salient; it does not reveal the level, horizon, or distribution of the searcher's belief.

Models can separate expectations from inflation premia

$$BE_t^{(n)} = E_t^P[\pi] + IRP_t + LP_t.$$

No-arbitrage and factor models combine yields, swaps, surveys, and inflation data to decompose a breakeven into a physical expectation, inflation-risk premium, and liquidity premium.

The decomposition is useful because it is model-based—and fragile for the same reason.

Central banks use measures to monitor anchors and test communication

- Compare long-run expectations across households, professionals, and markets to detect credibility gaps.
- Condition official forecasts on survey and market paths.
- Embed information treatments in ECB-CES and BoE DMP to estimate how messages change expectations and decisions.

Financial markets turn expectation measures into positions and hedges

- Trade the rate path with OIS and futures after policy announcements.
- Hedge purchasing power with TIPS, ILGs, and HICPxT swaps.
- Extract policy surprises from narrow-window futures changes for event-study regressions.

> **The same price can be a hedge instrument and an expectations signal.**

Corporations use expected rates and inflation in real decisions

- OIS and SOFR futures inform funding costs for variable-rate debt and planned issuance.
- Expected inflation enters wage bargaining and multi-year supply pricing.
- Long-run real-rate expectations determine investment hurdle rates.

> **Expectation measurement matters because contracts and investment are written before outcomes arrive.**

Researchers use observed beliefs to relax model-imposed expectations

- Put survey expectations directly into structural equations.
- Use cross-sectional disagreement to test sticky information, rational inattention, and information-friction models.
- Use survey RCTs to estimate belief-to-spending, hiring, investment, and price-setting chains.

Survey forecasts contain independent predictive information

> Michigan and Philadelphia SPF forecasts outperform nominal-yield-curve and time-series benchmarks for one-year US inflation in the evidence cited by the source.

Validation test: compare RMSE and out-of-sample predictive content against AR, VAR, and term-structure benchmarks.

> **Surveys are not dominated by financial prices; they contain a distinct signal.**

Stated plans predict costly behaviour

In the New Zealand firm evidence reported in the source, 97% of managers intending to change prices did so, while 98% stating no change made none. BoE DMP results are described as similar.

> Test realised prices, purchases, hiring, or investment against prior stated expectations and plans.

> **Survey answers are not cheap talk when they forecast real decisions.**

External records and RCTs complete the validation chain

Administrative alignment

Link survey microdata to scanner, bank-account, HR, investment, or text records.

Randomised causal chain

Assign information, verify that beliefs change, then test whether spending, hiring, investment, or prices change.

Validation requires beliefs and behaviour to move in sequence

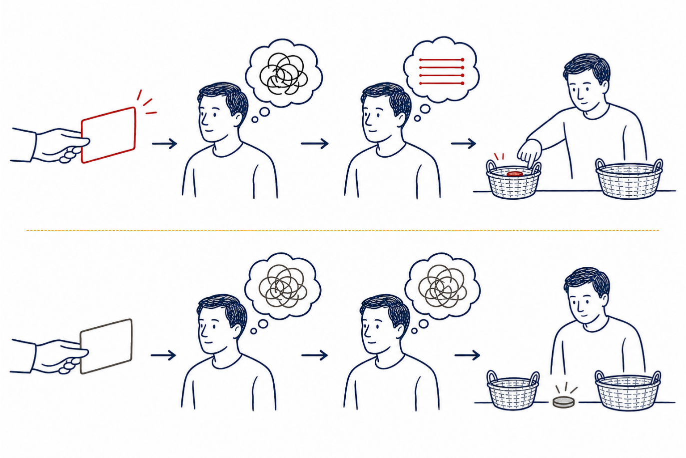

Editorial marker illustration generated for this lecture.

02B

Futures and asset prices

Prices, paths, surprises

Fed funds futures turn an overnight rate into a monthly contract

Instrument
Cash-settled CME future

→

Settlement object
Monthly average EFFR

→

Price convention
$F_{t,m} = 100 - \bar{r}^{Q}_{t,m}$

Trading began in 1988, formerly on CBOT. Final settlement uses the arithmetic average of daily EFFR values across the contract month—not a single day's rate and not the midpoint of the FOMC target range.

> **A contract price is interpretable only after we know the meeting date.**

The same contract serves expectations, hedging, and identification

Market expectations

A real-time reading of where traders price the Fed's near-term policy path.

Interest-rate hedging

Banks, asset managers, and corporations hedge short-term borrowing costs.

Academic identification

Kuttner (2001) uses event-window changes to isolate the unanticipated policy component.

Exchange	Settlement	Notional	Delivery	Minimum tick
CME Group	Cash only	USD 5 million	Monthly	1/2 bp

EFFR is the volume-weighted median transaction rate

$$\sum_{i=1}^{k-1} v_i < \frac{1}{2}\sum_{i=1}^{n} v_i \leq \sum_{i=1}^{k} v_i.$$

Order overnight federal-funds transactions from the lowest to the highest rate. Accumulate their volumes. The EFFR is r_k, the first transaction rate at which cumulative volume reaches half of total volume.

Because EFFR is the futures settlement object, its construction comes before contract interpretation.

A USD 50bn volume threshold selects 5.33%

Trade	Rate	Volume	Cumulative volume
1	5.30%	USD 10bn	USD 10bn
2	5.32%	USD 15bn	USD 25bn
3	**5.33%**	USD 30bn	**USD 55bn**
4	5.34%	USD 20bn	USD 75bn
5	5.36%	USD 25bn	USD 100bn

Total volume is USD 100bn, so the halfway point is USD 50bn. The third transaction is the first crossing: EFFR = 5.33%.

Before the median rule, outliers affected the fixing more strongly

$$\mathrm{EFFR}_{\text{pre-2016}} = \frac{\sum_{i=1}^{n} r_i v_i}{\sum_{i=1}^{n} v_i}.$$

The source contrasts the current volume-weighted median with the earlier volume-weighted mean. The median is less sensitive to a small number of unusually priced trades.

> The date label and historical wording should be checked against the New York Fed methodology before publication.

EFFR tracks the target range in the recent sample

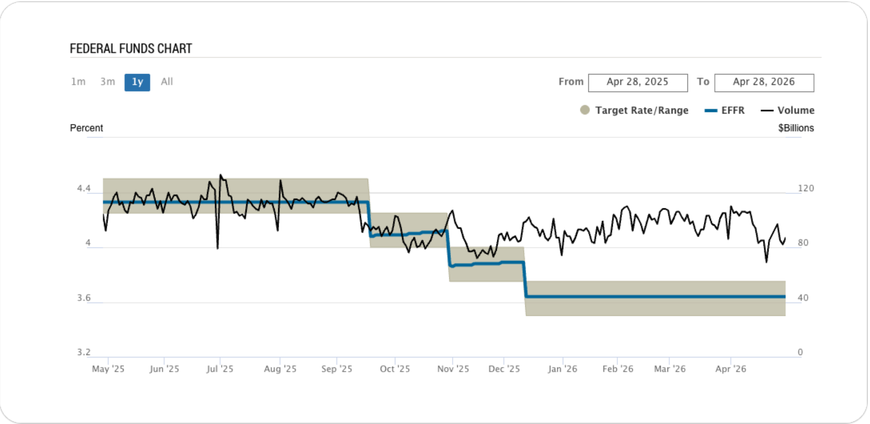

Federal Reserve Bank of New York, Effective Federal Funds Rate.

The long history reveals regime and operating-framework changes

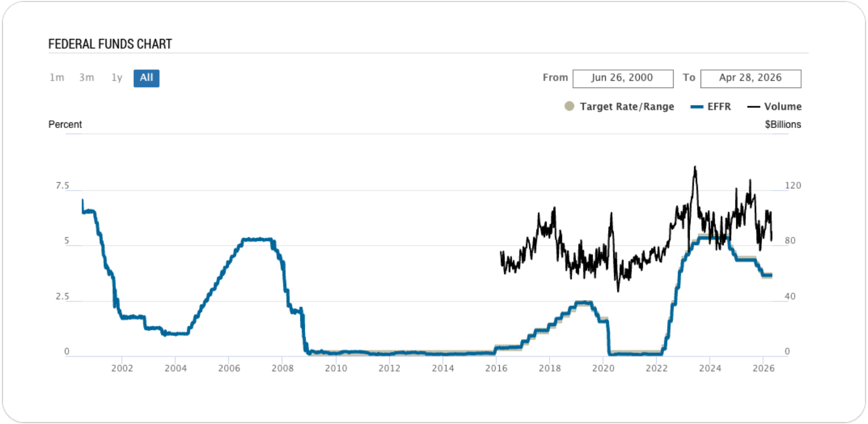

Federal Reserve Bank of New York, Effective Federal Funds Rate.

The IMM quote reverses the direction of the implied rate

$$F_{t,m} = 100 - R_{t,m}.$$

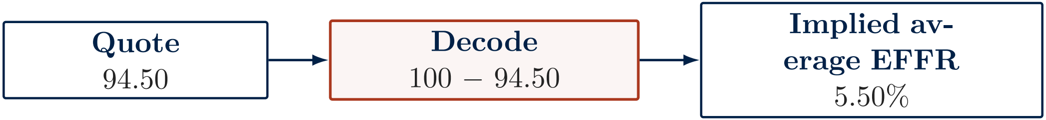

$F \uparrow \Rightarrow R \downarrow$: expecting a rate cut means buying the future.

Tick values translate quote precision into dollar exposure

Minimum fluctuation

0.5 basis point

The contract has roughly one month of annualised rate exposure.

Dollar values reported in the source

USD 20.835 per half-bp

One full basis point is USD 41.67.

> Tick value $= (\text{USD } 5{,}000{,}000 \times 0.00005)/12 = \text{USD } 20.835$.

Consecutive contract months trace a market-implied path

May 5.50% ← **Jun** 5.35% ← **Jul** 5.20% ← **Sep** 4.95% ← **Dec** 4.70%

Adjacent months are the raw material for recovering meeting-level expectations. Each month must still be decoded using its own calendar and any FOMC meeting contained within it.

> **A curve of monthly averages is not yet a curve of meeting-day target rates.**

Final settlement averages one EFFR input for every calendar day

Daily inputs
$r_1, r_2, \ldots, r_N$

$\rightarrow$

Arithmetic average
$\bar{r}_m = N^{-1} \sum_{i=1}^{N} r_i$

$\rightarrow$

Final price
$100 - \bar{r}_m$

Before expiry, the traded rate is approximately $F_{t,m} \approx 100 - \mathbb{E}_t^Q[\bar{r}_m]$. At expiry, realised EFFR inputs replace the expectation mechanically.

Weekends inherit the previous business day's EFFR

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
EFFR	5.33	5.33	5.58	5.58	5.58	5.58	5.58
Status	published	published	published	published	published	carry Fri	carry Fri

- Business-day observations enter once.
- Weekends and holidays carry forward the preceding published fixing.
- A change before a weekend can therefore enter settlement several times.

The meeting date determines the post-meeting weight

$$\bar{r}_{\text{month}} = \frac{(d-1)r_{\text{pre}} + (N-d+1)r_{\text{post}}}{N}.$$

Meeting on day d of an N-day month. Days $1, \ldots, d-1$ use r_{pre}; days $d, \ldots, N$ use r_{post}.

Thus $\omega_{\text{post}} = (N-d+1)/N$ and $\Delta F = -\omega_{\text{post}} \Delta r$.

The same 25 bp decision produces a different futures move on a different date.

The contract averages days before and after the meeting

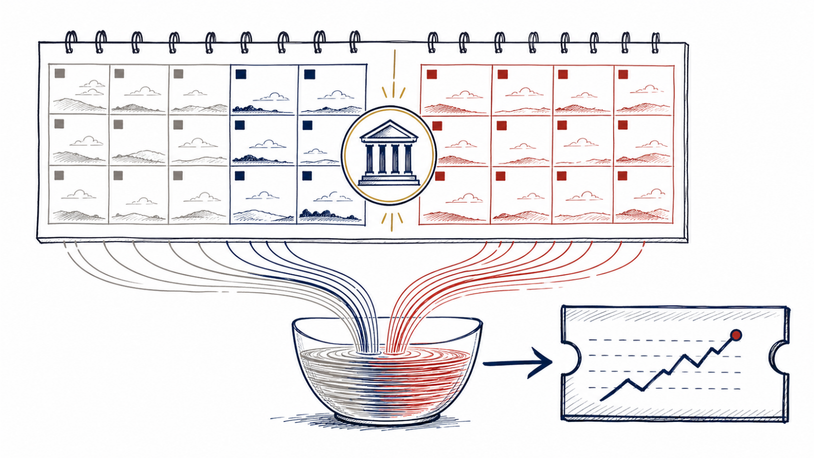

Editorial marker illustration generated for this lecture.

Early meetings load more heavily into the monthly average

Meeting day	Post-meeting days	Weight	Reading
1	30	100.0%	full month at new rate
5	26	86.7%	early-month meeting
15	16	53.3%	just over half
25	6	20.0%	late-month meeting
30	1	3.3%	single-day effect

Only a month with no FOMC meeting can be read directly as a constant expected rate.

Worked example: a 25 bp hike on day 15

Inputs

- $N = 30$ calendar days.
- $d = 15$ meeting day.
- $r_{\text{pre}} = 5.10\%$.
- $r_{\text{post}} = 5.35\%$.

Calendar split

days 1–14: 5.10%	days 15–30: 5.35%

The meeting day is counted in the post-meeting block, so 16 days carry the higher rate.

The calendar-weighted monthly rate is 5.233%

$$\bar{r}_{\text{month}} = \frac{14(5.10) + 16(5.35)}{30} = \frac{157.00}{30} \approx 5.233\%.$$

Block	Days $\times$ rate	Contribution
Old-rate block	14×5.10	71.40
New-rate block	16×5.35	85.60

The same hike settles differently when its date changes

$$\text{Settlement} = 100 - 5.233 = 94.767.$$

Meeting timing	$\bar{r}$	Settlement
Day 1	5.350%	94.650
Day 15	5.233%	94.767
Day 25	5.133%	94.867
No meeting	5.100%	94.900

FedWatch packages contract prices as target-range probabilities

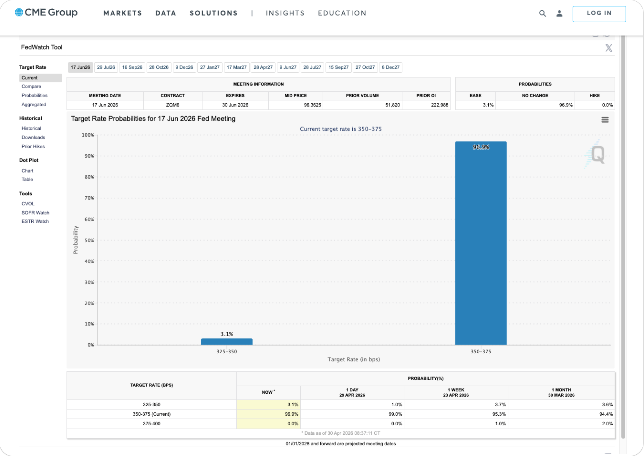

CME FedWatch Tool; screenshot embedded in ML2.

Meeting extraction is an inversion problem

1. **Read the pre-meeting rate.** Use f_0, the contract ending before the meeting, as an approximation to r_{pre}.
2. **Read the meeting-month average.** The contract containing the meeting gives f_1, a mix of old-rate and expected post-meeting days.
3. **Invert the average.** Solve for r_{post}, then compute $\Delta r = r_{\text{post}} - r_{\text{pre}}$.

Two contract readings plus a calendar recover a meeting-level expectation.

The meeting-month rate combines two policy regimes

$$f_1 = \frac{d-1}{N} r_{\text{pre}} + \frac{N-d+1}{N} r_{\text{post}}.$$

The pre-meeting rate is already largely observed. The only unknown component is the rate that will prevail after the FOMC decision.

Solve the monthly average for the expected post-meeting rate

$$r_{\text{post}} = \frac{N f_1 - (d-1) r_{\text{pre}}}{N - d + 1}.$$

The denominator is the number of post-meeting calendar days in the source's day-count convention.

Late-month meetings make the denominator small and the inversion noisy.

A 5.23% meeting-month rate implies roughly a 24 bp move

$$r_{\text{post}} = \frac{30(5.23) - 14(5.10)}{16} = \frac{85.5}{16} \approx 5.34\%.$$

$$\bar{r}_{\text{month}} = f_1 \quad \Longrightarrow \quad f_1 = r_{\text{expected constant}} \qquad \text{if there is no meeting.}$$

Against a 5.10% pre-meeting rate, 5.34% is an implied move of approximately +24 basis points.

A binary policy choice maps the expected move into a probability

If the only outcomes are a +25 bp hike and a hold, then

$$P_h^Q = \frac{\mathbb{E}_t^Q[\Delta r]}{25\text{ bp}} = \frac{r_{\text{post}} - r_{\text{pre}}}{25\text{ bp}}.$$

An implied move of +24 bp therefore maps to a 96% risk-neutral probability of a 25 bp hike.

The probability is conditional on the assumed outcome set.

Risk-neutral probabilities need not equal physical beliefs

$$P^{Q}_{\text{hike}} = P^{P}_{\text{hike}} + \lambda_t.$$

The additive wedge is a schematic reminder: futures prices live under Q. In calm periods the gap may be small; in stress it can be economically meaningful.

With three outcomes $(-25, 0, +25)$, one contract gives one equation for several probabilities. Options, additional prices, or structure are needed.

The Kuttner surprise scales a monthly move into an overnight shock

$$TS_t = \frac{N}{N-d+1}\Delta f_t = -\frac{N}{N-d+1}\Delta F_t.$$

F_t: futures price · $f_t = 100 - F_t$: futures-implied rate.

The gross-up factor converts the change in a monthly average into the implied surprise to the post-meeting overnight rate.

A surprise is the decision relative to the price just before it

$TS_t > 0$
Surprise tightening: more tightening or less easing than priced.

$TS_t < 0$
Surprise easing: more easing or less tightening than priced.

$TS_t = 0$
The action was fully anticipated.

A narrow announcement window holds scheduled macro news fixed, blocks policy feedback to within-window market moves, and limits movement in risk and liquidity premia.

> **The surprise is not the policy action itself.**

One futures surprise opened three empirical literatures

Study	Contribution reported in the source
Kuttner (2001)	Anticipated moves have little effect; surprises move bond yields.
Gürkaynak, Sack & Swanson	Separate target news from path or guidance news.
Nakamura & Steinsson	Use high-frequency surprises as instruments in structural macro models.

Forecast errors diagnose premia, bias, and regime change

$$FE_{t,h} = r_{t+h}^{EFFR} - f_{t,h}.$$

$$f_{t,h} = \mathbb{E}_t^P[r_{t+h}^{EFFR}] + \lambda_{t,h}.$$

If futures were unbiased physical forecasts with no wedge, $E_t[FE_{t,h}] = 0$. Persistent errors can instead reveal compensation for policy risk, forecast bias, or a structural break.

Predictive power fades as the horizon admits more news

Horizon	Source-reported MAE	Reading
1–3 months	5–10 bp	strong
3–6 months	15–30 bp	moderate
6–12 months	30–60 bp	weak
$>$ 12 months	$>$ 60 bp	very weak

FFF have tended to underpredict rises before tightening cycles and overpredict cuts before easing cycles. Beyond six months, premia and institutional breaks matter; the source flags the April 2023 Eurodollar-to-SOFR conversion.

Mechanics: quote, settle, then extract the meeting rate

$$F_t = 100 - \bar{r}_t^Q.$$

$$\text{Settle} = 100 - \frac{1}{N}\sum_{i=1}^{N} r_i^{EFFR}.$$

$$r_{\text{post}} = \frac{N f_1 - (d-1) r_{\text{pre}}}{N - d + 1}.$$

Price convention, settlement rule, and calendar determine the mechanical object.

Economics: probability, surprise, then the market wedge

$$P^{Q}_{\text{hike}} = \frac{r_{\text{post}} - r_{\text{pre}}}{25\ \text{bp}}.$$

$$TS_t = \frac{N}{N-d+1}\Delta f_t = -\frac{N}{N-d+1}\Delta F_t.$$

$$f_{t,h} = \mathbb{E}^{P}_{t}[r_{t+h}] + \lambda_{t,h}.$$

Probabilities and surprises become economic beliefs only after the Q-measure wedge is interpreted.

02C

High-frequency identification

Timing as identification

Policy reacts to markets and macro-financial news

$$\Delta i_t = \beta \Delta r_t + \gamma z_t + \varepsilon_t.$$

Δi_t: policy news · Δr_t: asset-price or yield change · z_t: common macro-financial news.

ε_t is the monetary-policy innovation we would like to isolate. At daily or monthly frequency, policy can also respond to market conditions and news.

Asset prices respond to policy and the same news

$$\Delta r_t = \alpha \Delta i_t + \delta z_t + \eta_t.$$

α: causal asset-price response to policy news · η_t: asset-market innovation.

The desired arrow is from Δi_t to Δr_t, but both variables also share z_t, and market shocks can feed back into policy.

The low-frequency data are generated by a simultaneous system.

A low-frequency OLS coefficient is not the causal response

$$\widehat{\alpha}_{\text{OLS}} \neq \alpha, \qquad \text{Cov}(\Delta i_t, \eta_t) \neq 0.$$

Simultaneity
Markets affect policy while policy affects markets.

Omitted news
z_t drives policy expectations and prices together.

Anticipation
Prices often move before the announcement.

The structural system makes both directions explicit

$$\begin{pmatrix} 1 & -\beta \\ -\alpha & 1 \end{pmatrix} \begin{pmatrix} \Delta i_t \\ \Delta r_t \end{pmatrix} = \begin{pmatrix} \gamma \\ \delta \end{pmatrix} z_t + \begin{pmatrix} \varepsilon_t \\ \eta_t \end{pmatrix}.$$

α is the market response to policy; β is the policy response to market conditions. Neither equation can be read in isolation.

The policy regressor already contains the excluded shocks

$$\Delta i_t = \frac{1}{1-\alpha\beta}\left[\varepsilon_t + \beta\eta_t + (\gamma+\beta\delta)z_t\right].$$

- ε_t: desired monetary-policy news.
- $\beta\eta_t$: market shock fed back into policy.
- $(\gamma+\beta\delta)z_t$: common macro-financial news.

The probability limit exposes both bias channels

$$\operatorname{plim} \widehat{\alpha}_{\mathrm{OLS}} = \alpha + \frac{(1-\alpha\beta)\left[\beta\sigma_\eta^2 + \delta(\gamma+\beta\delta)\sigma_z^2\right]}{\sigma_\varepsilon^2 + \beta^2\sigma_\eta^2 + (\gamma+\beta\delta)^2\sigma_z^2}.$$

Under the source's independent-shock calculation, market feedback enters through $\beta\sigma_\eta^2$ and common news through $\delta(\gamma+\beta\delta)\sigma_z^2$.

The same estimand can be written as covariance over variance

$$\widehat{\alpha}_{OLS} = \frac{\mathrm{Cov}(\Delta i_t, \Delta r_t)}{\mathrm{Var}(\Delta i_t)}$$

$$= \frac{(\beta\delta + \gamma)(\alpha\gamma + \delta)\sigma_z^2 + \alpha\sigma_\varepsilon^2 + \beta\sigma_\eta^2}{(\beta\delta + \gamma)^2\sigma_z^2 + \sigma_\varepsilon^2 + \beta^2\sigma_\eta^2}.$$

Bias disappears only when the unwanted variation disappears

$$\widehat{\alpha}_{OLS} - \alpha = (1 - \alpha\beta)\frac{\beta\sigma_\eta^2 + \delta(\beta\delta + \gamma)\sigma_z^2}{\sigma_\varepsilon^2 + \beta^2\sigma_\eta^2 + (\beta\delta + \gamma)^2\sigma_z^2}.$$

Identification is a signal-to-noise problem, not a larger-sample problem.

Simultaneity bias needs policy feedback and market shocks

$$\beta \neq 0, \qquad \sigma_\eta^2 > 0.$$

When the central bank responds to financial conditions, the policy-rate regressor contains the market innovation from the outcome equation. The regression then reverses part of the causal arrow.

Common-news bias needs news that enters both equations

$$\sigma_z^2 > 0, \qquad \delta(\beta\delta + \gamma) \neq 0.$$

CPI, labour-market releases, growth news, earnings, or changes in risk appetite can revise policy expectations and asset prices at the same time.

A close-to-close move can be precisely measured and still not be policy news.

Event time aims to make policy news dominate

$$\sigma_\varepsilon^2 \gg \sigma_\eta^2, \qquad \sigma_\varepsilon^2 \gg \sigma_z^2.$$

A narrow FOMC window does not assume that market shocks and common news never occur. It selects minutes in which monetary-policy news is expected to dominate their variation.

In the limiting thought experiment, OLS approaches the causal response

$$\frac{\sigma_\varepsilon^2}{\sigma_\eta^2} \to \infty, \qquad \frac{\sigma_\varepsilon^2}{\sigma_z^2} \to \infty \quad \Longrightarrow \quad \widehat{\alpha}_{OLS} \to \alpha.$$

A full day is often too wide: scheduled releases, firm news, and risk-premium shifts can all arrive. High-frequency studies use timing to shrink those components.

The announcement clock creates a narrow candidate window

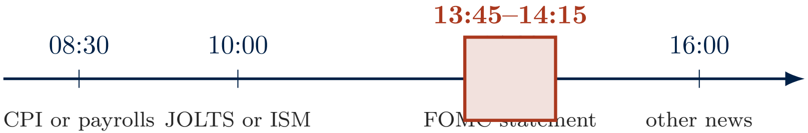

- No scheduled macro news inside $\mathcal{W}$.
- The policy decision cannot react to within-window price movements.
- Risk and liquidity premia are approximately stable inside the window.

A narrow window makes one filtered signal dominant

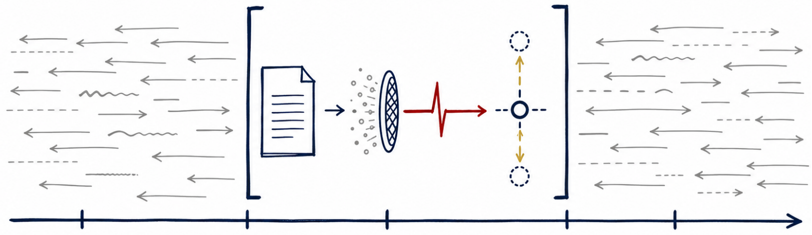

Editorial marker illustration generated for this lecture.

Measure both the asset and policy surprise over the same window

$$\Delta_{\mathcal{W}} A_{j,e} = A_{j,\tau^+} - A_{j,\tau^-}.$$

$$TS_e = \Delta_{\mathcal{W}} f_e.$$

e: announcement event · j: asset · τ^-, τ^+: just before and after the event.

The event-study regression still rests on exogeneity

$$\Delta_{\mathcal{W}} A_{j,e} = \alpha_j + \beta_j TS_e + u_{j,e}.$$

$$\mathrm{Cov}(TS_e, u_{j,e}) = 0.$$

Timing makes the covariance restriction more credible; it does not make it directly testable. Contemporaneous central-bank information effects can still violate the one-shock interpretation.

A narrow window excludes most of the trading day

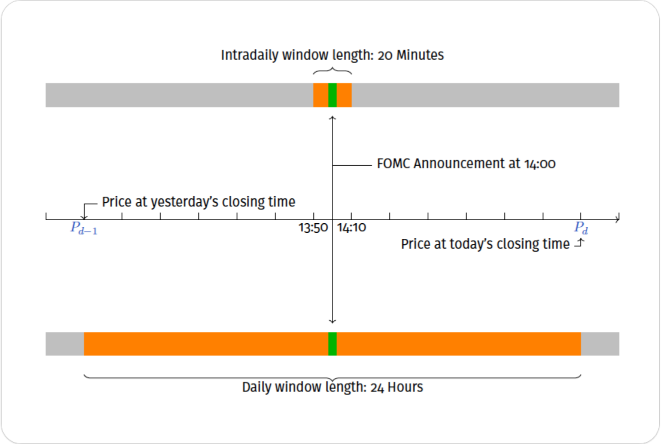

Event-window illustration embedded in ML2.

A daily window also contains the 08:30 macro release

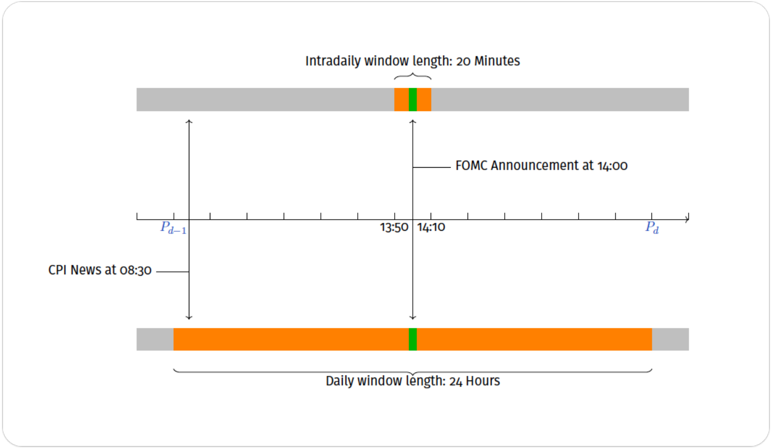

Event-window illustration embedded in ML2.

Adding 10:00 news compounds the daily-window confounding

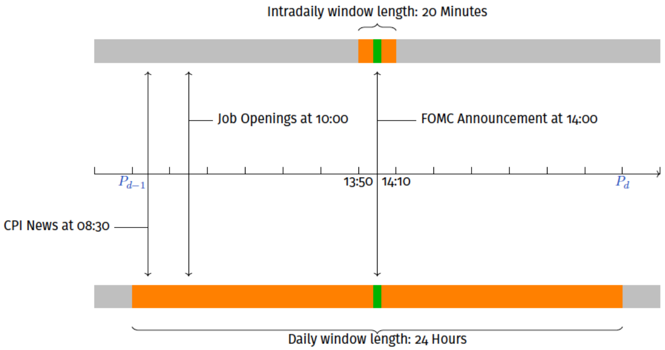

Event-window illustration embedded in ML2.

Late-day news makes close-to-close attribution still weaker

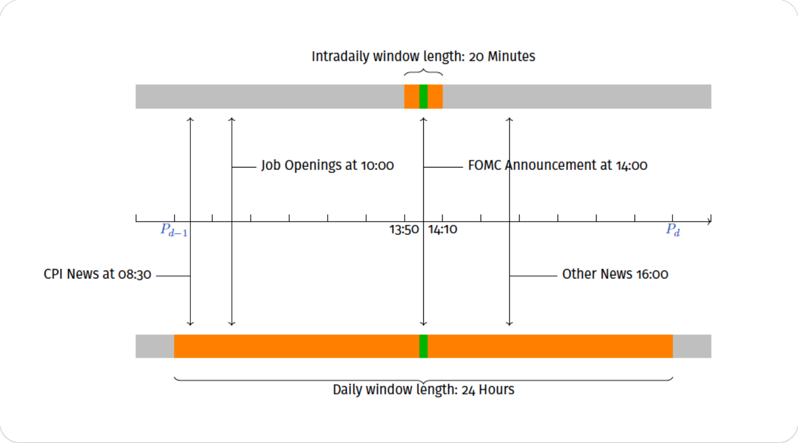

Event-window illustration embedded in ML2.

Intraday prices show exactly when an equity reprices

Amazon intraday chart embedded in ML2.

Euro-area OIS separates communication windows within one event

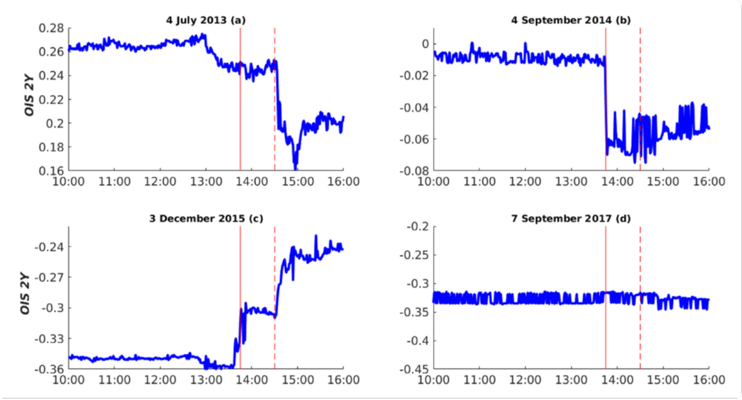

Euro-area intraday OIS evidence embedded in ML2.

Fed funds futures move sharply around monetary-policy news

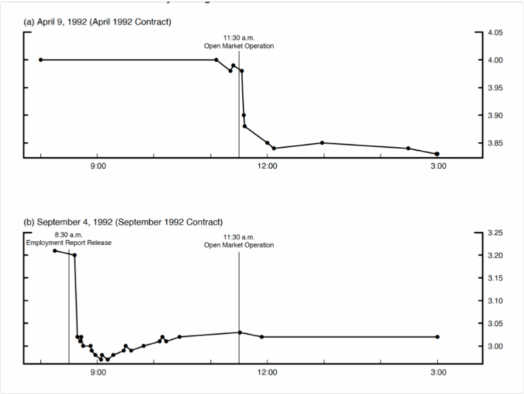

Intraday fed funds futures evidence embedded in ML2.

Intraday measurement brackets the exact announcement time

$$TS_d = \frac{D}{D-d}\left[-\left(ff_{\tau+\Delta\tau} - ff_{\tau-\Delta\tau}\right)\right].$$

τ: announcement time · $\tau - \Delta\tau, \tau + \Delta\tau$: quotes just before and after.

The narrow window offers stronger isolation but can amplify microstructure noise and requires intraday data.

Daily measurement is easier but admits every same-day shock

$$TS_d = \frac{D}{D-d}\left[-\left(ff_d - ff_{d-1}\right)\right].$$

The close-to-close change is available for more assets and is easy to replicate, but it includes scheduled releases, earnings, and general risk repricing on the announcement day.

Window choice trades contamination against market microstructure

Window	Gain	Cost
30 minutes	strongest isolation	noise; may miss delayed interpretation
2–3 hours	statement plus press conference	more overlapping market news
Daily	broad asset coverage; easy replication	weakest identification

Baseline: 30 minutes. Report robustness for 60 minutes, the press conference, and daily windows.

Window choice is an identification assumption, not a formatting choice.

Before the announcement, only part of the monthly average is uncertain

$$ff_{d_0-1,t} = \frac{d_0}{D_0} r_0 + \frac{D_0 - d_0}{D_0} E_{d_0-1}(r_1) + \eta_{d_0-1,t}.$$

Past calendar days contribute the known rate r_0. Remaining days contribute the expected post-meeting rate r_1, plus a market wedge.

After the announcement, the expected rate becomes realised

$$ff_{d_0,t} = \frac{d_0}{D_0} r_0 + \frac{D_0 - d_0}{D_0} r_1 + \eta_{d_0,t}.$$

The known component is unchanged. The announcement replaces $E_{d_0-1}(r_1)$ with r_1, while the event-window strategy assumes that the wedge moves little.

At mid-month, the announcement moves only half the contract

$$ff_{14,t} = \frac{15}{30} r_0 + \frac{15}{30} E_{14}(r_1) + \eta_{14}.$$

Half the month is already determined; half remains in expectation. A 25 bp surprise therefore moves the monthly futures rate by only 12.5 bp before scaling.

This section uses a different day-count convention from source slide 25. The inconsistency is preserved and flagged in day2.txt.

Subtracting pre from post cancels the known component

$$\Delta f_{d_0,t} = \frac{D_0 - d_0}{D_0} \left[r_1 - E_{d_0-1}(r_1) \right] + \Delta \eta_{d_0,t}.$$

The term $d_0 r_0 / D_0$ is present in both quotes and cannot move at the announcement. Only the remaining days transmit the policy news.

A stable event-window wedge isolates the overnight surprise

$$r_1 - E_{d_0-1}(r_1) = \frac{D_0}{D_0 - d_0}\Delta f_{d_0,t}.$$

The step uses $\Delta\eta_{d_0,t} \approx 0$. It converts the attenuated monthly-average change into the innovation to the post-meeting overnight rate.

The price convention flips the sign of the target surprise

$$TS_{d_0,t} = \frac{D_0}{D_0 - d_0}\Delta f_{d_0,t} = -\frac{D_0}{D_0 - d_0}\Delta F_{d_0,t}.$$

A surprise tightening raises the implied rate f, lowers the futures price F, and gives $TS > 0$. Surprise easing reverses every sign.

The scaling factor is the inverse post-meeting weight

$$\phi = \frac{D_0}{D_0 - d_0} = \frac{1}{\omega_{\text{post}}}.$$

The current-month contract averages across the whole month. Because only the remaining days respond to the decision, the observed change is mechanically attenuated.

Observed futures movement is the weighted overnight surprise

$$\Delta f = \omega_{\text{post}} \Delta r \quad \Longrightarrow \quad TS = \frac{\Delta f}{\omega_{\text{post}}} = \phi \Delta f.$$

Scaling recovers the overnight-rate unit from a monthly-average price.

A 12.5 bp monthly move can represent a 25 bp surprise

$$\phi = \frac{30}{30 - 15} = 2, \qquad TS = 2 \times 12.5 = 25 \text{ bp}.$$

At mid-month, only half the contract is exposed to the new rate. Grossing the observed 12.5 bp move up by two recovers the 25 bp innovation.

Late in the month, the current contract amplifies noise

$$\phi = \frac{31}{31 - 30} = 31.$$

One basis point of quote noise or wedge movement becomes a measured 31 bp target surprise after the mechanical gross-up.

The next-month contract avoids the late-month gross-up

$$TS_{d_0,t} = -\Delta F^1_{d_0,t} = -\left(F^1_{d_0,t} - F^1_{d_0-1,t}\right).$$

The following month's contract is fully exposed to the new target rate, so no current-month scaling is needed.

> Switching contracts fixes mechanical amplification, not liquidity, stale quotes, or policy-risk premia.

Kuttner regresses the event-day yield change on the surprise

$$R_{i,d} = \alpha + \beta_1 TS_d + \varepsilon_{i,d}.$$

$R_{i,d}$: yield change on asset i · TS_d: futures-based surprise · β_1: response to one basis point of surprise.

Efficient markets separate anticipated from unanticipated policy

$$\Delta i_d = \underbrace{TS_d}_{\text{unexpected}} + \underbrace{(\Delta i_d - TS_d)}_{\text{anticipated}}.$$

The anticipated component should already be embedded in yields before the meeting. Event-day price changes should therefore load on TS_d, not mechanically on the total target-rate change.

The empirical result shifted attention from actions to revisions

Component	Source-reported response	Interpretation
Anticipated move	approximately 0 bp	already priced
Short yield	approximately 1 bp	close to one-for-one
Long yield	0.4–0.6 bp	target news is not the full path

Gürkaynak, Sack and Swanson add target and path factors; longer yields load primarily on path news. Nakamura–Steinsson show that high-frequency surprises also shift forecasts, raising central-bank information effects.

> **What moves markets is the revision relative to expectations, not the announced action alone.**

The measurement problem has three layers

Signal
Which agent, variable, horizon, and moment?

→

Wedge
Survey design, risk, liquidity, text selection

→

Identification
Which timing isolates an unexpected revision?

1. Define the latent expectation before choosing its proxy.
2. Decode the proxy mechanically before giving it economic meaning.
3. Use timing and assumptions to separate a revision from confounding news.

No expectation series is self-interpreting.

Only the market branch feeds the futures-event estimator

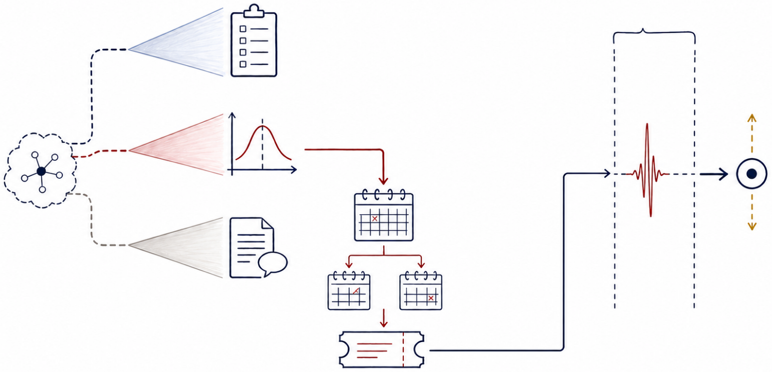

Editorial marker illustration generated for this lecture.

Signals · wedges · identification

1. Whose expectation did the measure capture?
2. Which wedge separates the observation from the belief?
3. Which timing assumption makes the revision causal?

Measure first; interpret second; identify third.

Sources: expectation measurement

- Lucas, R. (1976), "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Conference Series*.
- Binder, C. (2017), household inflation expectations and rounding, as named in the supplied lecture.
- Mankiw, N. G., Reis, R. and Wolfers, J. (2003), disagreement about inflation expectations, as named in the source.
- Ang, A., Bekaert, G. and Wei, M. (2007), "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" *JME*.
- Coibion, Gorodnichenko and coauthors; Bunn and coauthors; ECB-CES and BoE DMP validation and information-treatment evidence are retained as named in the HTML. Complete bibliographic details require the citation pass.
- Survey systems named in the lecture include Michigan, NY Fed SCE, Philadelphia SPF, Atlanta BIE, Cleveland SoFIE, BoE DMP, ECB-CES, and ECB SPF.

Sources: market instruments and data

- Federal Reserve Bank of New York, Effective Federal Funds Rate. **Methodology and data**
- CME Group, 30-Day Federal Funds futures. **Contract page**
- CME Group, FedWatch Tool. **Meeting probabilities**
- TIPS, ILGs, OAT/Bund breakevens; CPI, RPI, and HICPxT inflation swaps; SOFR, SONIA, and €STR OIS are the market measures named in the supplied lecture.
- Embedded EFFR, FedWatch, dot-plot, Bloomberg-style, and intraday charts are retained as documentary source assets rather than redrawn.

ML2 source slides 8–32; contract values and methodology dates require final publication-date verification.

Sources: policy surprises and identification

- Kuttner, K. (2001), "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market," *JME*.
- Gürkaynak, R., Sack, B. and Swanson, E., target and path factors, including the work identified as GSS (2005) in the source.
- Nakamura, E. and Steinsson, J., high-frequency external instruments and central-bank information effects, as named in the source.
- The simultaneous-system derivation, event-window restriction, and futures scaling equations are reconstructed directly from ML2 source slides 34–50.
- Source-reported yield responses and forecast-error ranges are preserved, but their table- or paper-level citations must be verified before publication.

Equations are native LaTeX; analytical timing diagrams are editable TikZ.

Sources: reconstruction note

- Content authority: `centralbanking_old/ml2.html`, all 51 top-level source slides.
- Visual authority: Central Banking Day 1 and `ml2026.sty`; palette, typography, title pages, dividers, footer, source notes, and density conventions are reused.
- All 77 source display-equation occurrences are re-typeset and logged in `day2.txt`; repeated summary equations remain repeated.
- Source images are kept in `assets/source`; presentation-ready copies are in `assets/normalized`.
- Six reviewed marker illustrations are integrated as dedicated conceptual frames. Briefs and generation records are in `assets/marker/illustration-requests.txt` and `assets/marker/illustration-prompts.txt`.
- Citation gaps, notation conflicts, and calendar-count inconsistencies are preserved and recorded rather than silently repaired.

Conversion principle: preserve first; improve only after a separate content review.