

2026-07-29 — the committee, simulated 19 times



Snapshot 07-15 • July Beige Book • released 2026-07-15, meeting minus 14 days

N = 19 runs

identical public inputs

gemini-3.5-flash

generated 2026-07-22

forward-looking • unscored


advance call locked 2026-07-20

Information cutoff 2026-07-15. The run sits at the end of a ten-day July window in which three documents landed: the Federal Reserve's Monetary Policy Report, published once a year and only in July; the Senate testimony Chair Warsh delivered alongside it; and the July Beige Book, released on the cutoff date itself. Nothing dated after that day is in the inputs.

Across 19 independent simulations of the 2026-07-29 FOMC on identical public inputs, the committee does not converge: the call splits 32% hike / 68% hold. No single run is quotable as a forecast.



The dispersion is **concentrated, not diffuse**: 5 members want a hike in every run, 2 never do, and **5 swing votes** decide the meeting.



01 The decision — direction, size and target range

Every one of the 19 runs opened from the same prevailing target range, **3.50% - 3.75%**, under a **neutral** regime anchor inherited from the last meeting — 19 of 19 runs carried that anchor, so nothing below reflects a regime the engine could not settle on. The runs produced **2 distinct outcomes** and no others.

OUTCOME	SIZE	TARGET RANGE	RUNS	SHARE
HOLD	0 bp	3.50% - 3.75%	13	68%
HIKE	+25 bp	3.75% - 4.00%	6	32%

There is no third scenario to hedge: no run stepped off the Committee's standard increment, and none moved twice. The hedge is therefore binary and both of its legs are named above — which is the part a probability alone does not give you. **The two branches are not equally comfortable, either.** The median dissent count differs by 3 votes between them, so one of these outcomes would be carried by a near-united Committee and the other over a bloc of objections. The headline statistics above give each branch its own figure for that reason: pooled across a divided arm, a median is a number no branch produced.

02 Member-level distribution — who are the decisive swing votes?

19 ballots per member. Each of the 19 re-runs is a complete meeting, so every member votes in every one: a row below is 19 readings of the same member, on the same date, from the same documents. It is that member's own distribution under identical conditions — not a voting record, and not 19 different meetings.

Read a row for its spread, not its total. A member who is 19 for 19 on one side is one the run-to-run variance never touched; a split row is a position the same inputs can argue either way, and those rows are where the meeting is decided. Position is read from each member's **own preferred rate**, not from Assent/Dissent — assent is relative to whatever the Chair tabled, so a member backing a hike in one run and a hold in the next looks identical in the vote column while having changed their mind.

MEMBER	HIKE	HOLD	CUT	HAWKISH	DISTRIBUTION	DISSENT	TRAIT	TYPE
Beth M. Hammack	19	0	0	100%		11%	+0.40	stable
Kevin M. Warsh	19	0	0	100%		0%	+0.55	stable
Lorie K. Logan	19	0	0	100%		21%	+0.50	stable
Michelle W. Bowman	19	0	0	100%		58%	+0.70	stable
Neel T. Kashkari	19	0	0	100%		26%	+0.30	stable
Christopher J. Waller	15	4	0	79%		16%	+0.10	SWING
Philip N. Jefferson	7	12	0	37%		11%	+0.00	SWING

MEMBER	HIKE	HOLD	CUT	HAWKISH	DISTRIBUTION	DISSENT	TRAIT	TYPE
Jerome H. Powell	3	16	0	16%		26%	-0.10	SWING
Anna L. Paulson	2	17	0	10%		32%	+0.00	SWING
John C. Williams	1	18	0	5%		42%	-0.20	SWING
Lisa D. Cook	0	18	1	0%		42%	-0.30	varies
Michael S. Barr	0	19	0	0%		42%	-0.20	stable

The **Dissent** column is a different question from position, and the answer is not the same. Assent is relative to whatever the Chair tabled, so a member who wants a hike in every run still assents whenever the Chair tables one — and dissents when the Chair tables a hold. Read the two columns together and the column tells you **who breaks ranks, and in which direction the Committee would have to go for them to do it**. That is the internal-disagreement signal the statement's tone follows, and it is invisible in a headline probability.

The personas are driving behaviour. Rank correlation between each member's simulated hawkishness and their configured trait is $\rho = 0.95$. Had this been low, the committee would be decorative and a member-level reading of these results would carry no information.

03 The rate path — where the disagreement actually sits

The dot plot across the 19 runs, split by what the meeting decided. **Read the branch columns, not the pooled median.** Pooling a divided arm produces a number that is nobody's projection: the median of the whole sample is simply the larger branch's, and on its own it reads as "no move all year" while 6 of the runs moved.

HORIZON	HIKE (6)	HOLD (13)	BRANCH GAP	ALL RUNS
2026	3.875	3.625	25 bp	3.625
2027	3.625	3.500	12.5 bp	3.500
2028	3.375	3.250	12.5 bp	3.250
longer run	3.125	3.125	0 bp	3.125

The dispersion in the path is the decision, not noise. Each branch is internally consistent: a run that moves carries that move to year-end, and a run that holds ends the year where it started. So the 2026 dot is 25 bp apart between the two branches — one step, exactly the step this meeting is deciding — and that gap decays with horizon until it is 0 bp at the longer run, where both branches sit at 3.125.

The desk reading is that **this meeting sets the level for about eighteen months and moves the destination not at all**. Whichever branch lands, the projected path from here is flat-to-lower after this year: neither branch projects a second move in the same direction. A pooled median hides both halves of that.

1 of 19 runs contradict themselves here — they raise the range and still end the year at the range they started from, pricing a reversal they never deliberated. Reported rather than dropped: it is one of the few places this engine can be caught being internally inconsistent, and it is the kind of thing that should shrink as the projection module improves.

04 Surprise risk — what the curve is not carrying

A distribution is only useful against another distribution. Below is the market-implied probability as of 2026-07-15 from CME FedWatch, against what these runs produced.

DIRECTION	MARKET-IMPLIED	SIMULATION	GAP
HIKE	15%	32%	+17 pp
HOLD	85%	68%	-17 pp

The gap is the surprise risk, and it is +17 pp on the hike. That is the part of the distribution priced into nothing: if the Committee moves, it moves against a curve carrying 15%. The engine is not a price model and does not try to be one — it is a second opinion built from the meeting rather than from the quote, and the two disagreeing by this much is itself the signal.

05 Why HIKE, when most runs said HOLD

HIKE +25 bp

Target range **3.75%–4.00%** · locked 2026-07-20 and unchanged since · traced to the deliberations of 12 simulated members.

The count that matters is **19 re-runs of the same meeting** — not 19 members, and not 19 scenarios. It is the minimum repetition this simulator needs before dispersion can be read as signal rather than noise. On byte-identical public inputs the meeting did not come out the same twice: **6 runs ended in a hike, 13 in a hold**.

That ratio is not the finding. **The dispersion is concentrated, not diffuse**: 5 members hiked in every run, 2 in none, and the 5 swing votes decided the meeting — with rank correlation between simulated hawkishness and configured persona at $p = 0.95$. These members are behaving like themselves, not like noise. A futures curve can price the outcome; it cannot tell you who moved, and it cannot tell you what happens if the Beige Book lands differently — there is no market in a hypothetical.

We publish **HIKE** as the headline direction against our own 68% hold, for two reasons. First, **backtests show this simulator misses turns at inflection points** — its failure mode is staying with the prevailing regime one meeting too long, which is precisely the error a modal reading would make here. Second, on 2026-07-15 **CME FedWatch put a hike near 15% while this simulation put it at 32%**. That gap is the part of the distribution the market was not carrying, and it earned weight.

That last call was ours, not the model's. The engine reports a distribution; deciding to publish against its own mode is interpretation, and interpretation is still human work — for now. The distribution is printed here unaltered, so a reader can reject the overlay and keep the numbers.

On 2026-07-29 the real committee sits **once**. 19 simulated meetings against the one that counts — the score gets published either way.

06 What ships with this snapshot

Every figure on this page is reproducible from the artefacts filed beside it. Same inputs, same settings, same result — so a number here can be quoted in your own research with its provenance intact.

- **19 complete simulated meetings** — the full deliberation transcript of each, with every member's argument, the Chair's proposal and the roll-call.
- **A drafted policy statement and meeting minutes** per run, in the Committee's own register — the tone and forward guidance read ahead of the release, not after it.
- **A dot plot per run**, aggregated into the path table above.
- `report_data.json` — every number on this page, language-agnostic, so nothing here has to be re-derived by hand to be checked.

07 Method & caveats

- The headline call is an editorial overlay on the run distribution, not the engine's modal output. Both are on this page and are never merged: the distribution is reported as the engine produced it, and the published direction is stated separately, with its reasons.
- **This is not a projection meeting.** The Committee publishes a Summary of Economic Projections four times a year and this meeting is not one of them, so no dot plot will be released on the day. The path above is the engine's own forced projection — useful as a read on where the simulated committee thinks the rate is going, and not comparable against anything the Fed will publish this month.
- Every figure here comes from N independent runs on byte-identical inputs; dispersion is the engine's own run-to-run variance, reported rather than averaged away.
- Member position is taken from each member's preferred rate (projection), not Assent/Dissent, which is relative to the Chair's tabled motion.
- Ballots within a run are NOT independent — members deliberate together. Effects are therefore tested with the run as the unit (per-run swing count, Mann-Whitney), never by pooling individual ballots, which inflates significance.
- Swing/core grouping is derived from the same runs it describes, so it is post-hoc; treat member-level effects as leads for a pre-registered follow-up rather than established results.
- This meeting has not occurred. No ground truth exists yet and nothing here is scored.

Simulation output — **not investment advice**. All inputs are public (FRED, published Federal Reserve documents). Dispersion is reported rather than averaged away; where a result is not statistically resolvable at this sample size, the report says so.