



EUROPEAN CENTRAL BANK

EUROSYSTEM

Discussion of FOMC *in Silico*: A Multi-Agent System for Monetary Policy Decision Modeling”

by *Sophia Kazinnik and Tara M. Sinclair*

Frankfurt am Main
23 March 2026

Fédéric Holm-Hadulla, Monetary Policy Department

The views expressed here are those of the speaker and do not necessarily reflect the views of the ECB



Motivation

Central bankers decide on monetary policy not just by casting a vote; decisions are usually preceded by in-depth deliberation in the monetary policy committee

Research question

Does this process of deliberation matter? In particular, does it produce policy rate decisions that deviate from those in a rules-based voting process?

Approach

- Build a **dual-track simulation**: LLM agents – standing in for FOMC members – deliberate in natural language (behavioral track) alongside a Bayesian voting model (rational track)
- Isolate the “**deliberation wedge**” as the gap between the outcomes in the respective tracks as to: rate decisions, dispersion between preferred rates, and dissents
- “Stress test” the results under two scenarios: **political pressure** on the Chair and a **negative revision in labor market data**

Scenario results (July 2025, FFR at 4.25–4.50%)

- 1. Baseline:** Committee converges unanimously on 4.42%. Deliberation introduces **hawkish bias** vs. benchmark. No dissents in either track.
- 2. Political pressure:** Chair's authority reduced; career-driven dovish shifts; views **polarize**; influence reallocates to likely successors.
- 3. Downward labor market revision:** Layered on political pressure. Final FFR drops to **4.30%**. Shocks compound / disagreement amplifies.

Validation, 215 meetings over 2000–2026:

93%

of meetings within 25 bps


MAE = **9.1 bp**; 54% exact

LLM reduces MAE by 29% vs. rational benchmark

(accuracy much lower in unconditional validation exercise)

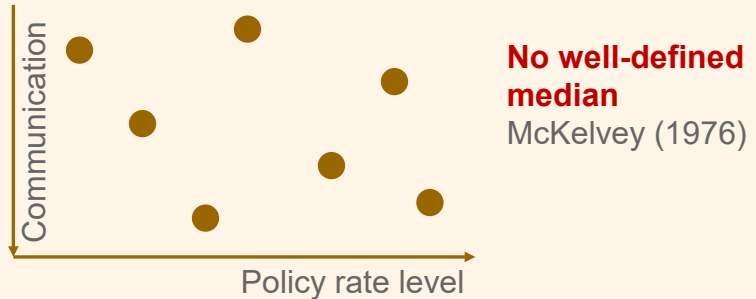
Dimensionality of monetary policy decisions

Current paper considers single dimension

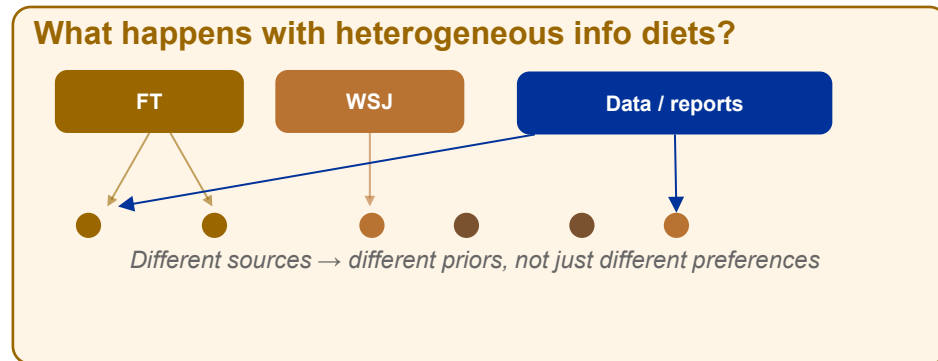
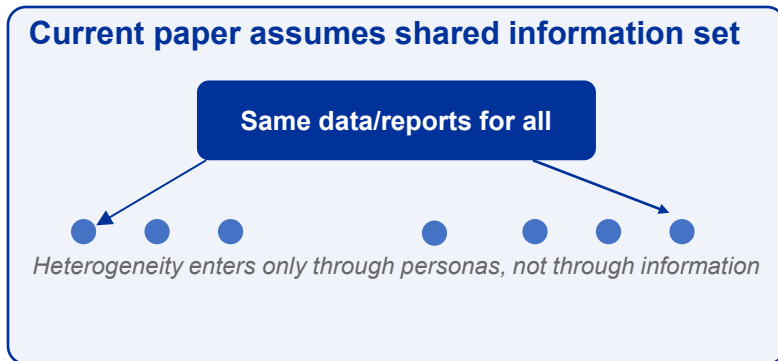


Median voter theorem holds, voting produces a unique Condorcet winner

Many CB decisions have multiple dimensions



- Rates are often **just one dimension of the decision problem**; others include communication, increment size (25 bp norm vs. 50/75 bps), forward guidance, QE.
- LLM agents already deliberate over multiple dimensions in natural language, but **rational benchmark relies on the Median Voter Theorem** to aggregate preferences.
- How **would deliberation wedge change if both tracks were harmonized** on the dimensionality of the decision problem?



- Growing evidence that the **source of information shapes beliefs**, even among professionals (Cookson, Engelberg & Mullins, 2023, RFS).
- How would dispersion / deliberation dynamics change **if different members consumed different info sets?**

Background

How the simulation works in detail

