In the Fed's Mind*

Ali Kakhbod UC Berkeley

Amir Kermani UC Berkeley & NBER

Bernardo Maciel UC Berkeley

This version: December 2025

Abstract

We introduce a novel text-based methodology to measure the Federal Reserve's real-time attribution of inflation to supply and demand factors. Using large language models (LLMs), we analyze the full corpus of FOMC meeting minutes from 1936 to 2025 in two steps. First, we extract statements citing causes of inflation. Second, we classify each as supply- and demand-driven. This produces a high-frequency panel tracking the Fed's evolving diagnosis of inflation. We find that inflation responds more strongly to perceived supply pressure, while interest rates respond more aggressively to demand-driven inflation. Embedding our measures into an augmented Taylor rule shows that this policy asymmetry becomes especially pronounced in the post-Volcker era. Perceived demand pressures predict higher risk-neutral yields, while perceived supply pressures predict increases in term premia, consistent with supply shocks increasing uncertainty about the future path of interest rates.

Keywords: Inflation, Monetary policy, Text data, Communication, Generative AI.
JEL classification codes: C10, E31, E52, E58

*First version: June 2025. We thank Michael Bauer, Jonathan Hazell, Emi Nakamura, Ricardo Reis, Jón Steinsson, and seminar participants at the Berkeley Macro Lunch Seminar, LSE Money-Macro Seminar, and UCL for very helpful comments. Ali Kakhbod is with the University of California, Berkeley, Haas School of Business, email: akakhbod@berkeley.edu. Amir Kermani is with the University of California, Berkeley, Haas School of Business, and NBER, email: kermani@berkeley.edu. Bernardo Maciel is with the University of California, Berkeley, Haas School of Business, email: bernardo_maciel@berkeley.edu.

Introduction

The sharp rise in inflation following the Covid-19 pandemic reignited a long-standing debate in macroeconomics: what are the underlying causes of inflation, and how should monetary policy respond? In particular, the distinction between supply-driven and demand-driven inflation has taken center stage. This distinction is not merely academic. In standard macroeconomic models, supply and demand shocks have markedly different implications for optimal monetary policy. A central bank that reacts aggressively to supply-driven inflation risks deepening a slowdown, while failing to respond to demand-driven inflation risks entrenching price pressures. Understanding how policymakers interpret the sources of inflation in real time is thus essential for evaluating and predicting monetary policy decisions.

Empirically, however, distinguishing between supply- and demand-driven inflation in real time is notoriously difficult. Traditional approaches usually rely on structural macroeconomic models. While valuable, these techniques typically require strong identifying assumptions and often rely on data that were not available to policymakers at the time of decision-making. Moreover, they reveal how inflation can be decomposed from an econometrician's perspective—not necessarily how central banks themselves viewed the shocks as they unfolded.

In this paper we build on the narrative approach pioneered by Friedman and Schwartz (1963) and expanded by Romer and Romer (1989, 2023) and introduce a novel text-based methodology to systematically analyze Federal Open Market Committee (FOMC) meeting records from 1936 to 2025, creating a high-frequency panel of the Federal Reserve's real-time attribution of inflation to supply (or cost-push) and demand components. Using a two-stage extraction and classification procedure grounded in macroeconomic sign restrictions, we construct imbalance measures that summarize the Fed's internal diagnostic narratives. We first document substantial time variation in the Fed's assessment of the relative importance of demand and supply forces in explaining inflation, with supply-related factors playing a particularly prominent role in the late 1970s and during the post-Covid inflation surge. We then show a pronounced policy asymmetry: interest rates respond more aggressively to perceived demand-driven inflation than to supply pressures, especially in the post-Volcker era. Finally, we show that the Fed's diagnosis is reflected in financial markets, with demand imbalances primarily raising risk-neutral yields, while supply pressures predominantly increase bond and equity premia and decrease equity prices.

We begin our analysis by leveraging recent advances in large language models (LLMs)

to extract and classify economic content from the full corpus of FOMC meeting. These records document the discussions underlying monetary policy decisions and provide a direct window into policymakers' real-time interpretations of macroeconomic conditions. Our approach proceeds in two steps. First, we use LLMs to identify statements that explicitly cite reasons for inflation or disinflation. Second, we classify each reason as demand-driven or supply-driven using a theory-based criterion: whether the factor is expected to move inflation and output in the same or opposite directions. We then aggregate these classifications into meeting-level "imbalance" measures that capture the normalized net share of inflationary versus disinflationary remarks within each category.

We show that these measures closely track well-known historical inflation narratives. They highlight supply-driven pressures during the oil shocks of the 1970s, a surge in disinflationary demand forces during the Great Recession, and a sharp post-Covid transition: early meetings in 2020 emphasize disinflationary risks, while 2021–2022 are dominated by inflationary supply factors associated with global supply chain disruptions and energy price increases following Russia's invasion of Ukraine.

To frame the empirical analysis, we next develop a simple New Keynesian model with demand and supply shocks that maps the Federal Reserve's perceived diagnosis of inflation into monetary policy, realized inflation, and asset prices. Under discretionary optimal policy, the Fed fully offsets demand disturbances but only partially accommodates supply disturbances because of the inflation–output trade-off. This asymmetry carries directly into asset prices. Demand shocks primarily affect nominal yields through the risk-neutral channel, reflecting changes in the expected path of policy rates. supply shocks, by contrast, reduce output and consumption, raising the marginal utility of consumption in adverse states and inducing a positive covariance between nominal bond payoffs and the stochastic discount factor. As a result, supply shocks raise term premia rather than expected short rates, while also increasing equity risk premia. The model delivers three clear and testable predictions: the Federal Reserve responds more aggressively to perceived demand shocks than to supply shocks; realized inflation is more strongly driven by supply disturbances; and asset prices reflect this asymmetry, with demand shocks moving risk-neutral yields and supply shocks raising term premia and equity risk premia.

Motivated by the model's predictions, we estimate an augmented Taylor rule and find a clear asymmetry in monetary policy responses: the Federal Reserve reacts more aggressively to perceived demand-driven inflation while partially looking through supply pressures. This differentiation across inflation sources emerges during the Volcker era and remains a defining feature of U.S. monetary policy in the decades that follow. Consistent with this asymmetric response, we find that demand-driven inflation pressures dissipate

relatively quickly, whereas supply pressures have more persistent effects on realized inflation. As a result, observed inflation outcomes are driven predominantly by supply-side disturbances.

We next take the model's asset-pricing predictions to the data. Using Adrian, Crump, and Moench (2013) ACM yield decompositions around FOMC announcements, we find that demand imbalances primarily raise the risk-neutral component of yields, consistent with a tighter expected policy path, whereas supply imbalances primarily increase term premia. In high-frequency windows, supply imbalances are associated with lower equity prices and higher volatility, while demand imbalances elicit a much weaker equity response. These patterns align with the model's asset-pricing implications: supply shocks cause nominal bonds to comove with equities and therefore command a premium, whereas demand shocks do not. We also find that standard pre-meeting macroeconomic indicators explain only 20 percent of the variation in the Fed's diagnostic narratives, indicating that FOMC discussions contain information beyond publicly available pre-meeting data. Moreover, controlling for macro news strengthens the association between our demand and supply imbalance measures and movements in risk-neutral rates and term premia.

Contribution to the Literature. Our paper contributes to four main strands of the literature. First, our paper relates to the empirical literature that decomposes inflation into demand and supply components. One strand estimates structural macroeconomic models, and uses historical shock decompositions to attribute inflation movements to aggregate demand, productivity and cost-push shocks (e.g., Smets and Wouters (2007)). More recently, a growing literature exploits input–output matrices and production networks to disentangle aggregate demand forces from sectoral supply constraints and bottlenecks, quantifying how shocks propagate across industries and contribute to inflation dynamics (e.g., Baqaee and Farhi (2022), di Giovanni et al. (2023)). More closely related to our approach, another strand decomposes inflation into demand and supply drivers using sign restrictions. Eickmeier and Hofmann (2022) propose a quarterly structural factor model identified with sign restrictions and estimated on a large set of inflation and real activity indicators, while Shapiro (2024) also relies on sign restrictions but implements them in a disaggregated setting, using sectoral price and quantity information from the monthly Personal Consumption Expenditures index to construct demand- and supply-driven inflation measures. Relative to this literature, our approach does not rely on the structural assumptions of a specific model nor on data that are not necessarily available to FOMC participants at the time of policy meetings.

Second, our paper relates to the empirical literature on the estimation of the Federal

Reserve's monetary policy rule. Taylor (1993) originally showed that a simple linear rule not only fits the Federal Reserve's decisions during the late 1980s and early 1990s, a period in which the Fed successfully achieved inflation targeting, but also can serve as a useful normative benchmark for anchoring long-run inflation expectations and achieving the Fed's objectives. Subsequent work by Clarida, Galí, and Gertler (2000), Cogley and Sargent (2001, 2005), and Kim and Nelson (2006) documents significant time variation in the Fed's policy function, with a markedly stronger response of monetary policy to inflation since the Volcker era. Orphanides (2002, 2003) instead emphasize that real-time mismeasurement of macroeconomic variables, particularly the output gap, played a central role in the Great Inflation of the 1970s, rather than changes in the policy rule itself. More recently, Bauer, Pflueger, and Sunderam (2024) focuses on time variation in perceptions of monetary policy and shows that professional forecasters' expectations also imply substantial changes over time in how monetary policy is perceived.

We complement this literature by showing that, in addition to time variation in the Fed's policy reaction, changes in the sources of inflation have also contributed to differential policy responses to similar deviations of inflation from the Fed's target. In this sense, our results on the estimation of an augmented Taylor rule are most closely related to Hofmann, Manea, and Mojon (2024), who build on the demand and supply decomposition of Shapiro (2024). Our finding that the Fed responds differently to inflation depending on its underlying source is consistent with a key prediction of New Keynesian models, namely that optimal monetary policy may partially look through inflation driven by supply shocks (e.g., Clarida, Galí, and Gertler (2000), Guerrieri et al. (2021), and Beaudry, Carter, and Lahiri (2023)), particularly when policymakers are less concerned about their reputation (Nakamura, Riblier, and Steinsson, 2025).

Third, our paper is related to the growing literature that uses natural language processing and large language models to study monetary policy. Lucca and Trebbi (2009) compute "hawkishness" scores from FOMC statements, while Gorodnichenko, Pham, and Talavera (2023) and Curti and Kazinnik (2023) construct measures based on the tone and facial expressions of the Fed Chair during post-FOMC press conferences. Hansen and Kazinnik (2023) uses ChatGPT to parse and classify the policy stance of FOMC announcements and compares the results with human classifications. Using FOMC transcripts, Hansen, McMahon, and Prat (2017) studies the impact of transparency on internal FOMC discussions; Cieslak et al. (2023) measure inflation uncertainty and show that higher uncertainty is associated with a tighter policy stance; Cieslak and McMahon (2023) show that more hawkish stances expressed in FOMC meetings are associated with lower risk premia; and Shapiro and Wilson (2022) infer time-varying central bank objectives from

policymakers' statements. Aruoba and Drechsel (2024) show that the text of Federal Reserve staff documents contains information beyond numerical forecasts. We contribute to this literature by showing that FOMC internal discussions contain valuable information about FOMC members' and staff's real-time assessments of the sources of inflation, with important implications for monetary policy and asset prices.

Finally, our paper contributes to the literature on monetary policy shocks. A key challenge in this literature is that different shock constructions are often associated with distinct implications for asset prices and aggregate outcomes (Brennan et al. (2024)). Our results show that, consistent with the predictions of Cieslak and Pflueger (2023), asset price reactions to monetary policy announcements depend critically on the Federal Reserve's assessment of the underlying sources of inflationary pressures. In particular, supply-driven inflation shocks have a large impact on term premia, whereas demand-driven shocks primarily affect risk-neutral rates. Accordingly, variation in empirical findings across different monetary policy measures may partly reflect differences in how these measures load on demand- and supply-driven inflation pressures.

The remainder of the paper is organized as follows. Section 3 presents a simple model and the intuition for the results. Section Section 1 discusses the data. Section 2 describes our methodology and explains how we use a large language model to identify and count the Fed's perceived causes of inflation over time. Section 4 presents the empirical results. Section 5 provides a calibration of the model. Section 6 concludes.

1 Data

Our empirical analysis combines textual and macroeconomic data from different sources, described in this section. Appendix Table D.1 reports the start and end dates and the number of observations for the main sources, with all observations recorded at the FOMC-meeting date level.

Our baseline dataset is the full corpus of FOMC meeting minutes and transcripts, spanning 832 meetings from 1936 to 2025.[1] For our empirical tests, however, we focus on the modern policy regime that begins after the Treasury–Fed Accord of 1951—specifically, meetings from 1960 onward to align with the PCE core and headline inflation series. Moreover, restricting the estimation to the post-Accord period ensures that all observations

[1]The format of the minutes changes over time. Between 1936 and 1967, we use the "Historical Minutes" from the Fed website; the "Memoranda of Discussion" between June 1967 and March 1976; the transcripts between March 1976 and December 1992; and the modern "Minutes" from 1993 onwards. We do so because minutes are more informative about economic conditions after 1993 and less so before that. Using minutes for the most recent period also allows us to avoid the five year embargo for the transcripts.

come from a coherent monetary policy framework.[2]

We collapse every macro series to the FOMC-meeting frequency by assigning to each meeting the value observed for its calendar month. Therefore, all variables enter the analysis at the FOMC meeting level. For inflation data, we obtain monthly core and headline PCE inflation data from the FRED. The main series of interest rates is the monthly federal funds effective rate, also from the FRED. Finally, to account for periods in which the Federal Reserve operated near the zero lower bound, we use the shadow interest rate constructed by Wu and Xia (2016) based on the idea of interest rates as call options introduced by Black (1995).[3]

2 Methodology

2.1 Inflation attribution procedure

We build a measure of the Fed's own diagnosis of inflation by coding every FOMC minute from 1936 to June 2025. We use a large language model (LLM) to read through these minutes and identify what reasons were discussed by the Fed as possible causes of inflation. Unlike proxies based on ex-post price indices, our measure reflects what policymakers said—and therefore believed—when they set policy. The procedure is divided into two stages.

First Stage. The first stage identifies the reasons discussed by the Fed as causes of inflation at the time of the meeting. To do that, we feed each minute text into an LLM and ask it to list all reasons that FOMC participants explicitly cite in the text as structural causes or underlying factors driving inflation. Throughout the process, we use a structured-object API to programmatically populate each response field corresponding to the questions. This methodology turns thousands of pages of minutes into a clean panel that researchers can merge, rerun, and extend quickly and conveniently. We use OpenAI's GPT-5 model for this task.

In order to ensure robustness and the quality of the results, we explicitly ask the LLM not to include self-referential answers such as "increasing inflation" or "inflation concerns". We also require an explanation for the classification with direct reference and

[2]See Eichengreen and Garber (1991) for a detailed discussion of monetary policy under the pre-Accord regime, which operated as an implicit price-level target and, therefore, required pegged nominal yields.

[3]We define the zero-lower-bound period as the two spans when the Federal Reserve held the target federal funds rate in the 0–0.25 percent range: first from December 16, 2008 to December 15, 2015, and again from March 15, 2020 to March 16, 2022. We obtain the shadow rate series from the Atlanta Fed's website.

quotes to the original text. For each identified reason, the model records three fields: (i) a concise label ("reason"), (ii) a quotation supporting that label ("explanation"), and (iii) whether the participants judged the reason to push prices up ("I''), down ("D''), or gave no clear direction ("na''). In addition, it also records the speaker (whether an FOMC participant or staff member discussed that specific reason) and assigns it an importance score. The full prompt and additional details are described in Appendix A.1.

Second Stage. In the second stage, we feed each reason–explanation pair identified in the first stage back into the LLM. Next, we instruct the LLM to classify each reason as driven by demand, supply (i.e., cost-push), expectations, monetary policy, or indeterminate/unclear.[4]

More specifically, we provide the model with an objective benchmark, the theoretical sign restrictions implied by demand- and supply pressures, to decide whether each reason is demand- or supply-driven. A shift in the aggregate demand curve to the right increases prices and the output, while a shift in the aggregate supply curve to the left increases prices but decreases output.[5]

We also classify reasons into two additional categories: expectations and monetary policy. The expectations category captures statements about perceived longer-run changes relative to prior readings, while the monetary policy category isolates references to the policy stance as the channel influencing inflation.[6]

In this stage, the LLM's main task is to determine—based on the provided explanation and surrounding context—whether the pressure in question moves inflation and output in the same direction (demand) or in opposite directions (supply). Pressures that are ambiguous or tautological (e.g., "prices increased" without reference to underlying structural causes) are classified as "other." The prompt is fine-tuned to address edge cases that we found to be prone to misclassification in earlier iterations. The complete prompt and ad-

[4]Throughout the paper, we refer to supply and cost-push pressures interchangeably.

[5]One concern is that including the reason label, in addition to the explanation passage, may bias the model toward using its knowledge of historical events rather than the text itself. For example, labels such as "Covid pandemic disrupting supply chains" or "Energy prices going up as a result of the oil shock" may prompt the model to rely on prior knowledge of these episodes. For robustness, we repeat the second-stage classification using only the explanation, omitting the reason label. The resulting measures are highly correlated with those from our main specification. We nonetheless retain the baseline approach because a single passage often contains multiple reasons, and the label helps the LLM identify which specific reason we refer to in that instance. We discuss additional robustness tests against look-ahead bias in Appendix F.4.

[6]For expectations, we further restrict reasons to those that concern changes in long-run inflation expectations, excluding statements about near-term inflation or generic references to anchored inflation. This discipline sharply reduces the number of reasons classified as expectations, but omitting it leaves our main results largely unchanged. In Appendix F.3, we also consider a refinement that accounts for changes in how expectations are discussed over time, particularly before versus after the Great Inflation.

ditional details are provided in Appendix A.2.

The purpose of this two-step procedure is to ensure robustness of the results by clearly separating the reading and identification task of the first step from the inference and classification task of the second. The first step is a straightforward natural language processing exercise. The second relies on the model's basic reasoning ability: we are asking the LLM to extrapolate a causal link between the shock and its separate effects on output and inflation using a "Econ 101" model of the world. This separation minimizes hallucination and look-ahead bias since we do not feed the full text of the minutes into the second step. [7] The result is a panel of demand and supply reasons—each tagged as inflationary or disinflationary—that tracks the Fed's stated view of inflation's causes over almost nine decades. The output of the LLM conveniently provides structured data to the econometrician that can be read as a table for further analysis. The diagram in Appendix Figure C.1 summarizes the two-stage procedure.

Illustrative examples. To illustrate, we consider the discussion from September 16, 2008 FOMC meeting – the day after Lehman's collapse. As we would expect, contractionary demand reasons dominate the discussions of the meeting. The LLM identifies twelve reasons discussed by the participants in that ocasion. Nine of these are reasons contributing to a decrease in inflation. The last is a negative-supply reason associated with the decline in prices of energy and other commodities.

Consider the reason labeled "weakening labor market," with the minutes noting that "Participants 'generally thought that the outlook for inflation had improved, mainly reflecting ... the weakening of the labor market." In our framework, classification rests on the sign co-movement implied by standard models: a demand disturbance moves inflation and the output in the same direction, whereas a supply disturbance moves them in opposite directions. Applying this logic, weakening labor market is a contractionary demand shock: it shifts the IS curve left, lowering the output gap; via the New Keynesian Phillips curve, lower slack reduces inflation pressure, so both the output gap and inflation move down. This is the economic basis for coding the example as "demand".

[7] A general concern when using LLMs for measurement is look-ahead bias: because the model is trained on text written after most of the meetings in our sample, it may implicitly use ex-post information when interpreting historical passages. Our baseline design already mitigates this risk by separating (i) a first-stage extraction task from (ii) a second-stage classification task, and by feeding the classifier only the local reason and context passage rather than the full minutes. As an additional robustness check, we implement an "entity neutering" procedure following Engelberg et al. (2025): we iteratively mask and paraphrase each reason–context pair until an adversarial LLM can no longer infer the meeting date (within a two-week window), and then rerun the second-stage classification and downstream analysis on the neutered text. The resulting demand/supply series are highly correlated with the baseline measures, and our main regression conclusions are unchanged. See Appendix F.4 for more details.

The model records a brief, text-anchored rationale (the "scratch" field) explaining the sign assignment (e.g., "weaker consumption → output ↓ and inflation ↓"). The rationale serves two purpose: it improves consistency by enforcing theory-based sign rules and gives researchers a transparent trail to verify decisions, especially in edge cases.[8] Full examples of the model's output are provided in Appendix E.

The next excerpt of the LLM output is from the June 15, 2022 meeting, while the Covid pandemic was still ongoing. The reasons discussed in that meeting were mostly supply related and associated with the consequences of the Covid lockdowns and the Russian invasion of Ukraine. The reason highlighted is identified as "ongoing supply bottlenecks". It refers specifically to a discussion of ongoing supply bottlenecks and rising energy and commodity prices adding to inflation risks. The full output is also available in Appendix E. The "scratch" field maps a clear sequence: lockdowns disrupt supply chains, pushing prices up and output down. Following that logic, the LLM assigns a positive effect of the reason on inflation and a negative effect of the reason on output. Finally, and consistent with these classifications, it tags this reason as "Supply".[9]

2.2 Summary Statistics

After detailing the structure of the raw output of our methodology, we now examine the distribution of reason classifications. Appendix Table D.2 contains the most common reasons identified for demand, supply, and expectations. The table shows that some of the common reasons classified as demand are related to differences between output and potential, such as a tight labor market and economic slack. Common reasons related to supply mention cost-push factors, such as energy prices and productivity. Common reasons related to expectations describe both declines and increases in longer-term inflation expectations.

It is worth noting that reasons related to wage increases appear under both demand and supply. This is because our prompts direct the LLM to distinguish between demand- and supply-related wage increases. As a result, strike-induced or statutory wage hikes are classified as supply pressures, while broad labor-market strength is tagged as demand. This disambiguation is possible because, in the second stage, the model receives both the label and the accompanying passage in which that reason was discussed; thus, identical labels can be mapped to different classifications when the surrounding context differs.

[8] Human auditing of model outputs is a central recommendation in the text-analysis literature (Gentzkow, Kelly, and Taddy, 2019). Moreover, eliciting explicit chain-of-thought reasoning has been shown to substantially improve the quality of large language model inference (Wei et al., 2022).

[9] Appendix E also reports two additional examples: Example 3, a staff demand reason from the late 1960s, and Example 4, a staff supply reason from the 1970s oil shock.

Concretely, explanations that point to factors such as contracts, strikes, and minimum-wage changes produce a supply classification, while references to tight labor demand or hiring pressures lead to a demand classification.

The Committee's views within a meeting often offset one another: one member stresses a tight labor market pushing prices up, while another might cite weak consumer spending pulling them down. To better understand the distribution of the Fed's perception, we count inflationary and disinflationary remarks separately. Let $N_{x,t}^{(+)}$ and $N_{x,t}^{(-)}$ denote, respectively, the number of reasons on date t from category x that the text describes as raising or lowering inflation, where x can be demand, supply, or expectations. These raw counts summarize the Committee's perceived balance of inflationary and disinflationary perspectives for each category and in each meeting. Table 1 reports summary statistics for all six series and for their totals.

Table 1: Summary Statistics for Reason Effects

component	direction	count	mean	std	min	25%	50%	75%	max
Demand	−	763	0.747	1.093	0	0	0	1	6
	+	763	1.396	1.594	0	0	1	2	9
Supply	−	763	1.511	1.742	0	0	1	2	10
	+	763	3.021	2.708	0	1	2	5	21
Expectations	−	763	0.125	0.451	0	0	0	0	4
	+	763	0.140	0.441	0	0	0	0	4
Total		763	6.940	4.270	0	4	6	9	36

Notes: For each FOMC meeting, we count how many demand, supply, and expectations reasons the participants describe as inflationary, "+", or disinflationary, "−". The main text details how these variables are constructed. The table reports cross-meeting summary statistics for each of the six series and for their sum ("Total''). All figures are counts of reasons.

Across meetings, the median lists six tagged reasons. Two facts emerge. First, inflationary tags outnumber disinflationary tags. Second, supply motives occur roughly twice as often as demand motives, placing supply at the center of the discussion. This pattern reflects our classification rule: we tag a reason only when the passage explicitly links it to inflation. Discussions of demand—employment or overall conditions—are excluded unless the text connects them to inflation, even though standard models imply such links. This discipline depresses demand counts more than supply counts. As a robustness check, we repeat the analysis also including discussions of output as implicit reasons for inflation in the counts; the main conclusions are unchanged.

Figure C.2 plots, for each meeting, the total number of reasons we extract—pooling

demand, supply, and expectations, inflationary and disinflationary directions. Dots mark individual meetings; the black curve is an exponential moving average to capture longer trends. Between the 1960s and the 1990s, the number of reasons per meeting remained remarkably stable, altough with a large variance. After the 1990s, this number has increased slightly over time. This is despite the fact that the lenght of the modern minutes (used from 1993 to 2025) is much shorter than transcripts, memoranda of discussion, and most of the historical minutes used in the earlier sample.

Appendix Figure C.3 examines composition. A cell is filled whenever a demand-, supply-, or expectations-related reason is mentioned, yielding a time-series indicator of when each category appears. Inflationary demand reasons occur throughout the sample but become scarce around the Great Recession; disinflationary demand reasons concentrate in recessions. Inflationary supply reasons also appear across the full period, and more frequently during the 1970s oil shock, the energy-related shocks preceding the Great Recession, and the post-Covid period in 2021–2022. Disinflationary supply reasons arise mainly in the 2010s as energy prices fell, especially after 2015 during the oil glut.

2.3 Post-Accord Narratives

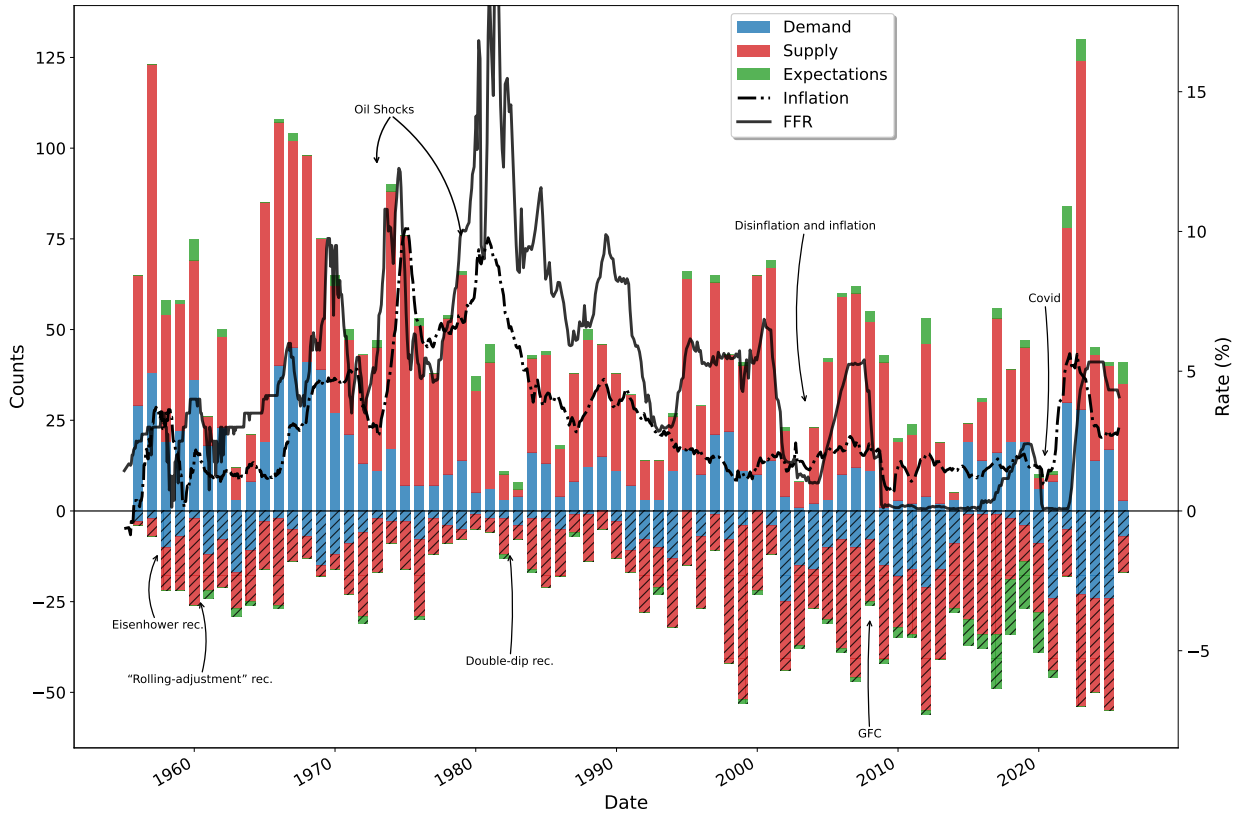
Figure 1 shows the annual totals of demand-, supply-, and expectations-related remarks (left axis). For context, the right axis overlays the federal funds rate and inflation.[10] The figure tells a story about some of the key monetary policy episodes in history. Next, we discuss some of these episodes in more detail, relating them to the reasons identified by our methodology.

Eisenhower recession (1957 – 58). In 1957, the Fed tightened to cool earlier price pressures as autos, steel, and housing slumped, compounded by an overseas liquidity squeeze and the Asian flu. Minutes record disinflationary forces—"weak sales", "rising joblessness", "inventory liquidation"—alongside inflationary ones—"defense spending", "easier credit", and restocking after the policy pivot. The Fed's quick rate cuts and fiscal stimulus ended the slide by mid-1958, but price pressures resurfaced in 1959 (Romer and Romer, 2002).

"Rolling-adjustment" recession (1960 – 61). Late 1959 inflation fears were driven by rapid inventory accumulation, strong car sales, and larger defense orders. A sharp turn

[10]For the ease of exposition, Appendix Figures C.4, C.5, and C.6 present snapshots of Figure 1 on three historical periods: the 1960s, the 1970s, and the Covid pandemic. Appendix Figure C.7 repeats the same plot, but using net imbalances (increasing minus decreasing counts).

Figure 1: Aggregated Reason Counts Over Time



Notes: This figure aggregates the demand, supply, and expectation reason counts annually, separately for reasons increasing (bars going up, full) and decreasing inflation (bars going down, hatched) in the left axis. Additionally, we present the Federal Funds Rate and inflation in the right axis. We use CPI inflation before 1960 and core inflation after that. See the main text for more details on the construction of the demand and supply measures from the minutes.

to inventory liquidation and unemployment above 6 percent flipped the narrative: the Committee now cited slack demand and balance-of-payments worries as disinflationary. Recovery began in mid-1961 as liquidation slowed and defense spending rose, yet persistent labor slack kept inflation subdued. The figure shows many inflationary mentions in 1960, then a marked shift toward disinflationary reasons through 1962.

Oil Shocks (1973, 1979). Throughout the 1970s, we see a relative increase in the presence of supply reasons discussed by the Committee. The 1973 and 1974 meetings mention "oil supply curtailment", the "fuel situation", "oil situation", "increased prices of commodities", "oil shortages" and "imported oil costs" as supply reasons contributing to an increase of inflation. Similarly, the 1979 meetings mention "higher prices of motor fuels", "energy price increases", "energy costs", while the 1980 meetings already account

for "energy price moderation".

Double-Dip Recession (1980, 1981 - 82). The graph shows mostly inflationary reasons in 1980, following by a drastic relative increase in disinflationary reasons in the following years. At the outset of 1980 the Committee blamed resurging prices on "energy costs,'' "wage growth'' that stayed at the "rapid pace'' of 1979, a "depreciating dollar", and a "widening foreign trade deficit''—all pushing inflation *up*. By March, however, the minutes pointed to "rising U.S. interest rates'' and a "new program of fiscal, energy, credit, and other measures'' that would push inflation *down*. Spring and summer entries discuss "weakening demand—declining retail sales,'' "rising unemployment,'' and "slow growth of money supply''—while also noting "energy price moderation.'' As the slump deepened in 1981-82 the discussion focused on a "strong dollar,'' "slower wage growth,'' "persistently high unemployment,'' and repeated "decline in real GNP,'' factors that sustained the disinflation despite temporary shocks such as the 1980 "surge in food prices."

Disinflation and Inflation (2000 - 2006). The graph shows a predominance of inflationary reasons in the 1990s, followed by a surge in disinflationary reasons up to 2003, when we see a reversal. At the turn of the decade, policymakers still saw disinflationary contributions from ''increasing productivity,'' expected to ''hold down increases in unit labor costs and prices,'' yet warned that ''sharp increases in energy prices'' and ''higher import prices'' together with ''taut labor markets'' were tilting risks upward.

By mid-2000 the list of cost-push worries had broadened to ''increasing medical costs'' and the ''rise in other import prices,'' while demand looked overheated thanks to ''strong consumer demand'' and ''extremely tight labor markets.'' The picture reversed in March, 2001: minutes stressed demand shortfalls—''declining consumer confidence,'' ''weak business investment,'' ''high inventory levels,'' and ''low resource utilization''—plus supply-side relief from ''falling energy prices'' and ''highly competitive conditions in most product markets.'' Indeed, Bernanke (2003) mentions a concern with disinflation at that point, with the Fed taking a more agressive expansionary monetary policy.

As the recovery firmed, the narrative flipped again: from 2004 on officials cited ''energy prices had risen steeply,'' ''unit labor costs had moved notably higher,'' ''tight resource utilization,'' and a ''recent decline in the dollar'' that risked ''putting upward pressure on import prices.'' By 2006 the Committee worried that "elevated energy prices'' together with ''continuing increases in resource utilization'' and signs that "inflation expectations had risen somewhat'' could entrench the rebound in inflation.

Post Covid (2020 - 2025). The 2020-2021 section of the graph is mostly characterized by disinflationary reasons, with a shift in 2022, when inflationary supply reasons resurfaced. Early in the crisis, the Committee stressed that "weaker demand and significantly lower oil prices are holding down consumer price inflation,"—a view that expressed concern about both demand- and supply-driven disinflation.

By mid-2021 the story flipped: the minutes note mostly cost-push pressures, such as "supply chain bottlenecks", "supply constraints", "input cost pressure", as the main contributors to inflation. In 2022, officials acknowledged that price pressures had become broad-based, mentioning both demand-related and supply-related reasons such as "demand-supply imbalances,"higher food and energy prices," and "Russia's war against Ukraine" as contributing to inflation.

By 2023, and after the interest rate liftoff, perceived demand and supply contributions fell, while the expectations component rose. This pattern is consistent with Reis (2022*a*), who attributes the 2021–2022 surge to misdiagnosed shocks—overstated slack and persistent supply constraints labeled transitory—and to a late recognition that expectations were drifting from target under a prolonged accommodative stance.

3 Model and Asset Pricing Implications

This section presents a parsimonious New Keynesian framework that links perceived demand and supply disturbances to inflation, monetary policy, and asset prices. All derivations are provided in Appendix B.

Inflation is not a sufficient statistic for optimal monetary policy and asset pricing responses because identical inflation readings can reflect different underlying shocks. In some meetings, yields move mainly through shifts in the expected short-rate path; in other meetings, they move largely through term premia, with long bonds behaving like risky assets and comoving with equities. Policy responses also vary: in some episodes the Committee tightens sharply, while in others inflation remains high and the policy path adjusts less than what the Taylor rule implies. These patterns fit poorly with a unidimensional view of inflation, but arise naturally when inflation is decomposed into different sources.

In the framework below, perceived demand and supply disturbances differ in their implications for equilibrium inflation, the policy-rate path, and the covariances that price nominal bonds and risky claims. Demand disturbances correspond to the "divine coincidence" case: under optimal discretionary policy, the rule offsets their effects on expected inflation and real activity, so their main pricing channel operates through the expected short-rate path. Supply disturbances instead generate a stabilization trade-off and

raise inflation while weakening real activity, which induces a positive covariance between nominal-bond payoffs and the pricing kernel. The model therefore maps the sign and magnitude of the perceived components into (i) the response of the expected short-rate path, (ii) the term premium in the yield decomposition, and (iii) the price response of the risky asset.

Macro Block. Let $x_t := y_t - z_t$ denote the output gap, π_t inflation, and i_t the nominal policy rate (all in deviations from long-run values, in logs). The system satisfies the New Keynesian IS and Phillips curves:

$$x_t = -\psi(i_t - \mathbb{E}_t \pi_{t+1}) + \mathbb{E}_t x_{t+1} + \kappa_d d_t, \tag{1}$$

$$\pi_t = \lambda x_t + \beta \mathbb{E}_t \pi_{t+1} + \phi_s s_t, \tag{2}$$

where d_t and s_t are demand and supply shocks with autoregressive laws of motion

$$d_t = \rho_d d_{t-1} + \varepsilon_{d,t}, \qquad s_t = \rho_s s_{t-1} + \varepsilon_{s,t}, \tag{3}$$

with ρ_d, $\rho_s \in (0,1)$, κ_d, $\phi_s \in (0,+\infty)$, and innovations $\varepsilon_{d,t}$ and $\varepsilon_{s,t}$ independent and mean zero. The policy instrument are short nominal interest rates (i_t) and expectations are common knowledge.

Under discretionary optimal policy with quadratic loss, the equilibrium policy rule is linear in the perceived components:

$$i_t = \mathbb{E}_t \pi_{t+1} + \Upsilon_s s_t + \Upsilon_d d_t, \tag{4}$$

where Υ_s, Υ_d depend only on structural parameters. The key implication is an asymmetry: demand disturbances are offset one-for-one in equilibrium, while supply disturbances generate a stabilization trade-off. Under standard parameter values,

$$0 < \Upsilon_s < \Upsilon_d, \tag{5}$$

so the real rates respond more to demand than to supply pressures.[11]

Asset Prices. We price nominal bonds and a risky claim using a standard stochastic discount factor implied by CRRA utility. In equilibrium under optimal policy, inflation and

[11] As we show later, the same holds for nominal rates under our calibrated parameters.

the output gap depend only on the perceived supply component,

$$\pi_t = k_\pi^s s_t, \qquad x_t = k_x^s s_t, \tag{6}$$

with

$$k_\pi^s = \frac{\alpha\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} > 0, \qquad k_x^s = -\frac{\lambda\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} < 0. \tag{7}$$

Importantly, demand disturbances are fully neutralized and do not enter (6).

Aggregate consumption equals output, $C_t = Y_t$, and $c_t = \log C_t$. With CRRA preferences, the log stochastic discount factor is given by

$$m_{t+1} = \log\beta - \gamma\,(c_{t+1} - c_t). \tag{8}$$

Using $y_t = x_t + z_t$ and the CGG interpretation that supply shocks lower output and potential output by the same amount, let $z_t = \bar{z} - z_s s_t$ (with $z_s > 0$). Output therefore falls with supply shocks: $y_t = \bar{z} + y_s s_t$ with $y_s < 0$. Hence,

$$m_{t+1} = \tilde{m}_t - \gamma\, y_s s_{t+1}, \tag{9}$$

where $\tilde{m}_t$ is a term known at time t.

We start by investigating the model implications for bonds. For the two-period nominal yield, the log-linear decomposition (following Cieslak and Pflueger (2023)) is

$$\begin{aligned} i_t^{(2)} = {} & \frac{1}{2}(i_t + \mathbb{E}_t i_{t+1}) + \frac{1}{2}\mathrm{Cov}_t\left(m_{t+1}, \pi_{t+1}^e\right) \\ & - \left[\frac{1}{2}\mathrm{Cov}_t\left(\pi_{t+1}, \pi_{t+1}^e\right) + \frac{1}{4}\mathrm{Var}_t\left(\pi_{t+1}^e\right)\right], \end{aligned} \tag{10}$$

where $\pi_{t+1}^e := \mathbb{E}_{t+1}\pi_{t+2}$. The first term is the risk-neutral expectations component; the covariance term captures the inflation risk premium. The last term is a Jensen's inequality adjustment term.

Using Equation (4) and the autoregressive dynamics for the demand and supply shocks, the risk-neutral component is linear in (s_t, d_t):

$$\frac{1}{2}(i_t + \mathbb{E}_t i_{t+1}) = \theta_s s_t + \theta_d d_t, \tag{11}$$

and under standard calibrations $\theta_d > \theta_s$, so demand shocks move the expected path of short rates more than supply shocks.

The inflation risk premium is driven by supply volatility. To see this, note that Equation (6) implies $\pi^e_{t+1} = k^s_\pi \rho_s s_{t+1}$. Combined with Equation (9), we then have

$$\text{Cov}_t\left(m_{t+1}, \pi^e_{t+1}\right) = -\gamma\, y_s\, k^s_\pi\, \rho_s\, \text{Var}_t(\varepsilon_{s,t+1}) \;>\; 0, \tag{12}$$

since $y_s < 0$. supply shocks make nominal bonds risky because they induce an endogenous rise in rates, so bond prices fall when consumption is low; in this sense, bonds behave like stocks. By contrast, demand-driven inflation carries no risk premium in this baseline because the endogenous policy response fully offsets it.

Finally, we consider the model implications for the equity market. Let $P^c_{1,t} = \mathbb{E}_t[M_{t+1}C_{t+1}]$ denote the price of a one-period consumption claim. A standard log-linear approximation implies

$$\mathbb{E}_t[r^c_{t+1} - r_t] + \frac{1}{2}\text{Var}_t(r^c_{t+1}) = -\text{Cov}_t(m_{t+1}, r^c_{t+1}) = \gamma y_s^2 Var_t(s_{t+1}) > 0. \tag{13}$$

Because output and consumption fall when s_t rises, the right-hand side is positive when supply shocks are present. Absent supply shocks, consumption is fully stabilized and the premium is muted.

4 The Fed's diagnostic of inflation and monetary policy

In this section, we first turn the minutes into meeting-level demand and supply imbalances —the net share of inflationary minus disinflationary reasons—normalized across categories. We then show that these imbalances are associated with innovations to inflation, with supply related imbalances tracking inflation more strongly.

Next, plugging the imbalances into a smoothed, augmented Taylor rule reveals a clear asymmetry: the policy rate responds strongly to demand-driven inflation and partially looks through supply pressures, with this pattern strongest in the period between Chair Volcker and the Great Financial Crisis. We then show that this pattern results in a persistent impact of supply imbalances on inflation and little to no relationship between past demand imbalances and inflation.

Finally, yield curve decompositions indicate that demand pressures raise risk-neutral yields, whereas supply pressures raise term premia. This finding is consistent with the New Keynesian stabilization trade-off: supply pressures force the central bank to balance inflation against economic activity, increasing uncertainty about the future policy path (and thus term premia), while demand pressures involve no such trade-off and primarily

lift the expected path of short-term interest rates.

We also test the predictability of the Fed's perception of demand and supply imbalances and find that pre-meeting measures of macroeconomic conditions explain only a small share of the Fed's diagnosis. Moreover, our constructed imbalance measures predict yield changes beyond those macroeconomic characteristics.

4.1 Empirical measure for the Fed's diagnostic

The previous section shows that our methodology broadly captures the economic discussions in some of the key historic monetary policy episodes. Next, we use the reasons identified previously to construct measures of the demand and supply proportional contributions to inflation.

As Figure C.2 showed, reason counts differ across meetings, and a single meeting can cite both inflationary and disinflationary forces within the same category. To address these two facts, we normalize the within-date counts and define the category-x imbalance, for $x \in \{d, s\}$, as

$$I_{x,t} := \frac{N_{x,t}^{(+)} - N_{x,t}^{(-)}}{\sum_{z\in\{d,s\}} \left| N_{z,t}^{(+)} - N_{z,t}^{(-)} \right|}. \tag{14}$$

where we denote demand and supply d and s. To limit noise from minor mentions, we restrict the sample to reasons with importance scores of at least five.

The normalization above corresponds to a L^1 normalization of the vector of net counts. Imbalances are always inside the interval $[-1, 1]$, with $\sum_{z\in\{d,s\}} |I_{z,t}| = 1$ by construction. A positive value means upward forces in category x dominate downward ones; a negative value means the reverse. Dividing by the absolute sum of imbalances across all categories reflects a relative contribution which takes into account other reasons discussed in that meeting. Appendix F.1 describes two alternative normalizations. The main results are robust to these choices.

4.2 Imbalances and inflation

Our first task is to test how the perceived supply and demand imbalances covary with realized inflation. Specifically, we estimate the following specification:

$$\pi_t = \alpha_0 + \beta_d I_{d,t} + \beta_s I_{s,t} + \rho \pi_{t-2} + u_t \tag{15}$$

where π_t is either core or headline inflation. As an initial validation, we expect the imbalances to co-move positively with inflation.

Table 2: Regressions of inflation on imbalances

	Core	*Headline*
I_d	0.064^{**}	0.123^{***}
	(0.026)	(0.045)
I_s	0.131^{***}	0.261^{***}
	(0.031)	(0.050)
π_{t-2}	0.981^{***}	0.961^{***}
	(0.017)	(0.020)
Obs.	630	630
R^2	0.979	0.957

Notes: The dependent variables are monthly levels of headline and core inflation, expressed as year-over-year percentage changes, from the FRED. Demand and supply imbalances are defined as in Equation (14). Regressions are estimated by OLS with Newey-West standard errors with 8 lags (in parentheses). *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table 2 reports the results. Demand and supply imbalances both covary positively with inflation. Coefficients are smaller for core than for headline, consistent with core excluding food and energy. Discussion of inflationary reasons intensifies in more inflationary periods—especially for supply.[12]

4.3 Augmented Taylor rule

Under the standard Taylor rule, the policy rate should rise more than one-for-one with inflation, regardless of its source. Table 2 shows that demand and supply imbalances both covary positively with inflation, with a slightly larger coefficient for supply. We next estimate an augmented Taylor in which we allow demand and supply imbalances to have distinct impact on monetary policy.

Let us consider the standard Taylor rule with interest rate smoothing:

$$i_t = \beta_0 + \rho i_{t-2} + (1-\rho)\left[\beta_\pi \pi_t + \varphi_x x_t\right] + \varepsilon_t. \tag{16}$$

Even though explanations for why such a smoothing exists differs in the first place, this

[12] Appendix Tables D.9 and D.14 report results under the alternative normalizations described in Appendix F.1. Results are very similar across these specifications.

is a widely accepted feature of monetary policy in the real world.[13] As before, i_t is the interest rate, π_t is inflation, x_t is the log output gap, and β_0 is a reduced-form intercept.[14] To handle zero lower bound periods, we set i_t to the shadow rate of Wu and Xia (2016) when available (1990-2022) and to the federal funds rate otherwise. We use the policy rate lagged two meetings (i_{t-2}) in the right-hand side of the equation above to approximate a quarterly lag used by most of the literature.

We define a Taylor rule that includes our demand and supply components:[15]

$$i_t = \beta_0 + \rho i_{t-2} + (1-\rho)\left[\varphi_\pi \pi_t + \varphi_x x_t + \varphi_d I_{d,t} + \varphi_s I_{s,t}\right] + \varepsilon_t. \tag{17}$$

Following the methodology of Carvalho, Nechio, and Tristão (2021), we can estimate the structural specification of Equations (16) and (17) by OLS in two steps. First, we define the auxiliary reduced-form equation:

$$i_t = \beta_0^{aux} + \rho i_{t-2} + \varphi_\pi^{aux} \pi_t + \varphi_x^{aux} x_t + \varphi_d^{aux} I_{d,t} + \varphi_s^{aux} I_{s,t} + \varepsilon_t. \tag{18}$$

We then estimate Equation (18) by OLS. Finally, we obtain the structural coefficients as

$$\widehat{\varphi}_{\text{var}} = \frac{\widehat{\varphi}_{\text{var}}^{\text{aux}}}{1-\widehat{\rho}}, \quad \text{for} \quad \text{var} = \pi, x, s, d.$$

We then compute standard errors using the delta method.

4.4 Taylor rule estimates under different monetary policy regimes

The history of monetary policy reveals substantial shifts in the Fed's view of the sources of inflation and its understanding of macroeconomics.[16] Clarida, Galí, and Gertler (2000) distinguish the post-Accord, pre-Volcker chairs (1951–1979) from the Volcker–Greenspan era (1979–2006). In the earlier regime, policy was accommodative: real rates fell as expected

[13] A number of papers have been written on this topic, and a comprehensive overview is outside the scope of this text. An early account by Goodfriend (1991) documents the preference of the Fed for interest rate smoothing. Sack and Wieland (2000) argue that interest rate smoothing might be optimal with forward-looking expectations or uncertainty about the data and the parameters of the model. Coibion and Gorodnichenko (2012) use forecasts from the Greenbook and the SPF and finds evidence that the interest rate inertia likely comes from smoothing, rather than from persistent shocks.

[14] Sack and Wieland (2000) and others make a distinction between this type of policy rule and the original Taylor rule presented by Taylor (1993). We don't make such a distinction, and for the purposes of our analysis, refer to all these variations of interest rate policy responses as "Taylor rules".

[15] Hofmann, Manea, and Mojon (2024) perform a similar exercise, augmenting the baseline specification with demand and supply components of inflation using the decomposition from Shapiro (2024).

[16] For the early U.S. experience, see Friedman and Schwartz (1963); for recent surveys, see Blinder (2022) and Bernanke (2022).

inflation rose. In the later regime, policy was proactive: real rates rose with expected inflation.

Our decomposition of inflation into demand and supply components allows us to revisit this debate and measure how the Fed's responses varied across these distinct periods. Table 3 re-estimates our augmented Taylor rule specification by subsample. In all specifications, the dependent variable is the shadow rate whenever it is available (1990-2022), and the federal funds rate otherwise.[17]

Table 3: Augmented Taylor Rule Estimates

	Full sample	*Pre-Volcker*	*Volcker–GFC*	*Post-GFC*
ϕ_π	1.37***	0.68**	1.81***	3.47*
	(0.34)	(0.27)	(0.19)	(2.05)
ϕ_x	1.04***	1.03***	0.68***	0.31
	(0.24)	(0.34)	(0.16)	(0.4)
ϕ_s	0.32	0.97	-0.13	-0.56
	(0.48)	(0.6)	(0.3)	(1.43)
ϕ_d	1.66*	0.12	1.74***	1.99
	(0.87)	(0.72)	(0.52)	(1.6)
ρ	0.87***	0.88***	0.62***	0.93***
	(0.03)	(0.04)	(0.07)	(0.04)
N	638	270	229	139
R^2	0.931	0.919	0.901	0.959

Notes: Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (17), for different samples: *Full sample* = 1960Q1–2025Q2; *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. φ_π is the response to π_t (core inflation), φ_x to x_t (log output gap); φ_d and φ_s to the demand and supply textual imbalances. ρ is the interest-rate smoothing parameter. Eq. (18) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

The first column presents the baseline results of our full sample from 1960 to 2025. The response to inflation is above unity; the incremental response to demand is positive and large, while the response to the supply component is small and statistically insignificant. This is consistent with the New Keynesian intuition of responding more to inflation when its perceived cause is demand rather than supply. Furthermore, the estimated ρ is positive and large, consistent with the well-documented persistence in policy rates.

[17] The results are qualitatively the same when we use the federal funds rate at all times.

The second column covers the pre-Volcker era—Chairs Martin (1951–70), Burns (1970–78), and Miller (1978–79).[18] The baseline coefficient on inflation is below unity, consistent with the accommodative stance documented by Clarida, Galí, and Gertler (2000). The imbalance coefficients are small and insignificant, suggesting the absence of a targeted policy response during this period.

The third column restricts the sample to the period between Volcker and the Great Financial Crisis and yields markedly different results. The baseline coefficient on inflation now exceeds one, indicating an aggressive anti-inflation stance. In addition, the demand imbalance coefficient is large and highly significant, while the supply coefficient is again small and insignificant.

The last column restricts the sample to the post-GFC period (2008–2025). The zero lower bound was binding for much of this period; consequently, persistence dominates the interest rate dynamics in this sample, and most estimates are not statistically significant. Nevertheless, the point estimate for demand remains positive, while that for supply is negative and smaller in magnitude. More importantly, as we show in the next section, imbalances still strongly co-move with the yield curve during this period, linking the Fed's perceptions of inflation's sources to market expectations about the future path of interest rates.[19]

A potential concern is attenuation of the supply coefficient because supply shocks raise contemporaneous inflation and affect the output gap—both included as controls. We therefore re-estimate the response of policy rates to perceived imbalances excluding core inflation and the log output gap. Table D.5 reports the results. The demand-imbalance coefficient exceeds the supply coefficient in all samples except the pre-Volcker period, consistent with a less differentiated policy stance before Volcker.

Across regimes, we find a clear shift in Fed behavior. Before 1979, the inflation coefficient is below one, and the policy rate does not differentiate between demand and supply assessments, consistent with the accommodative stance discussed in Clarida, Galí, and Gertler (2000). After Volcker's arrival, the inflation coefficient exceeds one, and the weight on demand-driven inflation is much larger than on supply pressures. These patterns indicate a departure from the standard rule and align with "targeted" Taylor rule evidence. The Fed moved from a passive, one-size-fits-all policy to an active rule that leans strongly against shifts in demand while partially looking through supply effects. After the Great

[18] The regression begins in 1960, when the FRED's core inflation series becomes available.

[19] Appendix Tables D.10 and D.15 report results under the alternative normalizations described in Appendix F.1. Results are very similar across these specifications. Appendix F.3 includes expectation imbalances in the estimated policy response. Additionally, Appendix F.2 repeats this exercise separately with reasons presented by FOMC members and staff.

Financial Crisis, unconventional tools became commonplace, and the policy rate became relatively less responsive to inflation and its perceived sources.

Persistence of demand and supply perceptions. Standard New Keynesian models imply that policy should stabilize demand-driven fluctuations but cannot fully offset supply shocks. If the Fed responds more to demand than to supply, lagged effects should decay faster for demand. We re-estimate Equation (15) with block-averaged imbalances—the mean of contemporaneous through three-meeting lags and the mean of four–seven-meeting lags —and control for inflation four and eight meetings earlier (about half a year and one year). Table 4 shows that inflation's dependence on imbalances fades faster for demand than for supply.

Table 4: Response of inflation to lagged imbalances

	Dependent variable: π_t			
	Core		*Headline*	
	(1)	(2)	(3)	(4)
$I_{d,t}$ to $I_{d,t-3}$	0.146	0.044	0.281*	0.225
	(0.098)	(0.217)	(0.161)	(0.294)
$I_{s,t}$ to $I_{s,t-3}$	0.469***	0.527***	0.831***	1.059***
	(0.101)	(0.157)	(0.148)	(0.212)
$I_{d,t-4}$ to $I_{d,t-7}$		0.095		0.007
		(0.162)		(0.232)
$I_{s,t-4}$ to $I_{s,t-7}$		0.637***		0.827***
		(0.171)		(0.216)
π_{t-4}	0.941***		0.893***	
	(0.031)		(0.035)	
π_{t-8}		0.844***		0.758***
		(0.053)		(0.058)
Obs.	632	628	632	628
R^2	0.946	0.855	0.906	0.791

Notes: The table reports OLS regressions of inflation π_t on demand and supply imbalances at the FOMC-meeting frequency. For each imbalance $(I_{d,t}, I_{s,t})$, we use block-averaged lags: the mean of contemporaneous through 3-meeting lags and, where included, the mean of 4–7 meeting lags. Columns (1) and (2) use core inflation, while columns (3) and (4) use headline inflation. Columns (1) and (3) include only t to $t-3$ for the imbalances and control for four lags of inflation (π_{t-1}–π_{t-4}), capturing short-run persistence. Columns (2) and (4) include both t to $t-3$ and $t-4$ to $t-7$ for the imbalances and instead control for π_{t-8} (roughly one year) to absorb longer-run persistence. Newey–West standard errors with 8 lags are reported in parentheses. See the main text for more details on the construction of the imbalances.

A natural question is whether this asymmetry reflects different persistence of the underlying shocks. One of the common explanations for weaker responses to supply is that

supply shocks are less persistent, so partially looking through supply-driven inflation is less costly. To test this, we estimate separate autoregressive models, using again averages of adjacent lags for each imbalance. Specifically, we regress contemporaneous imbalance on theree subsequent averages: averages of lags one to three, averages of lags four to six, and averages of lags seven to nine. Appendix Figure C.8 plots the estimated coefficients. Both imbalances exhibit similar persistence in our sample; demand and supply are equally as persistent in the Fed's discussions about the causes of inflation.

4.5 Yield curve responses and the Fed's diagnostic of inflation

The theoretical effect of demand and supply shocks on bond yields is ambiguous ex ante. For example, Pflueger (2025) shows that, in a model with external habits, the response of nominal yields to a demand shock can be either positive or negative depending on the calibration. By contrast, when monetary policy reacts more aggressively to supply shocks, interest rates move more with output, making bonds less attractive to risk-averse investors. In that case, bonds become more "stock-like," and the effect on nominal yields is large.

In this section, we examine how the Fed's assessment of the sources of inflation is reflected in Treasury yields and term premia. Movements in the yield curve reflect both the market's expectations of the future path of short rates and the compensation investors require for bearing interest-rate risk. We use the decomposition of Adrian, Crump, and Moench (2013), which separates Treasury yields into expectations and term-premium components.[20] Let $y_t^{(n)}$ and $tp_t^{(n)}$ denote, respectively, the fitted zero-coupon yield and the term premium at maturity n, measured two business days after FOMC meeting t. The corresponding risk-neutral yield $rn_t^{(n)}$ is then given by

$$rn_t^{(n)} = y_t^{(n)} - tp_t^{(n)}. \tag{19}$$

The term premium $tp_t^{(n)}$ captures the extra compensation for bearing interest-rate risk, while the risk-neutral yield $rn_t^{(n)}$ captures the implied path of future short rates under the risk-neutral measure after stripping out this risk compensation.

We regress risk-neutral yields and term premia of different maturities on the imbalances, controlling for core inflation, the log output gap, and the dependent variable lagged by two meetings, analogously to Equation (18). As such, these regressions capture the comovement between the Fed's assessment of the causes of inflation, measures by our imbalances, and the change in yields/premia between the current meeting and two meetings ago.

[20] Data from the Federal Reserve Bank of New York's website.

Table 5 presents the results using the entire sample, covering August, 1961 to June, 2025. The left panel shows the results for the risk neutral yields. We see that the demand imbalance is positively related to yields, with a front-loaded impact along the yield curve; by contrast, the effect of supply imbalances is much smaller. The right panel presents the effects on term premia and shows a very different picture. Demand imbalances are unrelated to term premia for any maturity, while supply imbalances are positively related to term premia along the curve, especially at the long-end.

Table 5: Changes in yields and term premia following the Fed's diagnostic

	Risk Neutral Yields				*Term Premia*			
	1y	2y	5y	10y	1y	2y	5y	10y
I_s	0.067	0.071*	0.066**	0.054**	0.031*	0.040*	0.043*	0.052*
	(0.043)	(0.040)	(0.033)	(0.026)	(0.016)	(0.020)	(0.025)	(0.031)
I_d	0.184**	0.183**	0.161**	0.127**	-0.010	-0.007	-0.000	-0.006
	(0.082)	(0.078)	(0.066)	(0.051)	(0.024)	(0.030)	(0.035)	(0.045)
π_t	0.137***	0.127***	0.104***	0.079***	0.013**	0.017**	0.020**	0.026***
	(0.042)	(0.035)	(0.027)	(0.020)	(0.006)	(0.007)	(0.008)	(0.010)
x_t	0.091***	0.079***	0.059***	0.043***	-0.010***	-0.016***	-0.021***	-0.019**
	(0.021)	(0.019)	(0.014)	(0.011)	(0.004)	(0.005)	(0.007)	(0.008)
$rn_{t-2}^{(n)}$	0.875***	0.877***	0.882***	0.886***				
	(0.028)	(0.025)	(0.023)	(0.021)				
$tp_{t-2}^{(n)}$					0.831***	0.874***	0.913***	0.922***
					(0.036)	(0.026)	(0.020)	(0.018)
Obs.	614	614	614	614	614	614	614	614
R^2	0.942	0.944	0.945	0.947	0.788	0.850	0.905	0.914

Notes: *Notes*: Each column reports OLS estimates of risk-neutral yields $rn_t^{(n)}$ or term premia $tp_t^{(n)}$ on demand and supply imbalances (I_d, I_s), core inflation (π_t), the log output gap (x_t), and the dependent variable lagged by two meetings, analogous to the auxiliary regression in Equation (18). Maturities are one, two, five, and ten years. Risk-neutral yields and term premia come from Adrian, Crump, and Moench (2013); yields and premia are measured two business days after FOMC meeting t. The sample spans August 1961–June 2025 at the FOMC-meeting frequency. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Appendix Table D.3 reports the same regressions for the post-GFC period (2007–2025). We observe the same asymmetry: demand imbalances load positively on the risk-neutral component, while supply imbalances load positively on term premia. Indeed, the estimated coefficient on term premia on supply imbalances are much larger and more significant for that subsample. Despite the federal funds rate's diminished role as a policy instrument, the Fed's perceived diagnostic of inflation remained systematically related to

yield-curve movements in this period.[21]

The asymmetric response to demand and supply imbalances is in line with the intuition from our model. Demand-driven inflation is associated with a stronger tightening path, which shows up as a large, front-loaded increase in risk-neutral yields and little movement in term premia. By contrast, supply imbalances generate a policy tradeoff between stabilizing inflation and real activity. This tradeoff makes the future policy path riskier and is reflected in higher term premia, especially at longer maturities, with a smaller effect on risk-neutral yields.

4.6 How predictable are the Fed's discussions?

In this section, we study to what extent the imbalances we extract from the minutes and transcripts are predictable by macroeconomic variables known prior to the announcements. We consider a set of pre-announcement macro and financial variables from Bauer and Swanson (2023): the latest nonfarm payroll surprise, year-over-year nonfarm payroll employment growth, three-month S&P 500 returns, the three-month change in the yield-curve slope, three-month changes in commodity prices (BCOM), and the implied skewness of the ten-year Treasury yield.[22]

Table 6 presents the results of regressions of the demand and supply imbalances we construct on this set of macro news predictors. A few interesting patterns stand out. First, the past performance of the S&P 500 positively predicts a demand imbalance, and negatively predicts supply imbalances. The change in commodity prices, as expected, is strongly related to the supply imbalance. Overall, both the change in nonfarm payroll and the nonfarmm payroll surprise are positively related to all measures. The R^2 is around 20%

[21] A concern is that movements in long yields may largely reflect repricing at short horizons. As a check, we repeat the analysis using Treasury forward rates implied by the Gürkaynak, Sack, and Wright (2007) yield-curve estimates (available from the Federal Reserve Board website). These series are not decomposed into risk-neutral and term-premium components, but they allow us to trace horizon-by-horizon movements in the term structure. We find that both demand and supply imbalances load significantly on forward-rate changes, with demand effects concentrated at shorter horizons and supply effects relatively more pronounced at longer horizons, which is consistent with the results above.

[22] The nonfarm payroll surprise (NFP Surprise) is the difference between the most recent nonfarm payroll release and the median expectation for that release from the Money Market Services survey. We standardize the surprise in the NFP to make it similar to the units of the other controls. Employment growth (NFP) is the log change in nonfarm payroll employment from one year earlier to the most recent FOMC announcement. S&P 500 returns are computed over the three months preceding the announcement. Yield Curve Slope is the change in the slope from three months prior to the announcement to one day before the announcement. Commodity prices (BCOM) are the log change in the Bloomberg Commodity Spot Price Index from three months before the announcement to the day before the announcement. Treasury Skewness is the implied skewness of the ten-year Treasury yield over the preceding month. See Bauer and Swanson (2023) for more details. The data is obtained from the San Francisco Fed's website.

Table 6: Imbalance Regressions on Macro News

	I_d	I_s
S&P500	1.180***	-1.346**
	(0.424)	(0.609)
BCOM	0.793*	3.351***
	(0.417)	(0.606)
NFP Surp.	0.020	0.039
	(0.024)	(0.025)
NFP	0.079***	0.077***
	(0.022)	(0.025)
Slope	-0.022	0.138
	(0.070)	(0.101)
TR Skew	-0.300**	0.191
	(0.146)	(0.232)
Obs.	284	284
R^2	0.222	0.206

Notes: The table reports the regressions of the demand and supply imbalances on the set of macro news predictors from Bauer and Swanson (2023). See the main text for the details on the variables. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

for demand and supply suggesting that, although recent macroeconomic conditions captured by this list of predictors are related to the Fed's assessment of the causes of inflation, there is substantial variation that remains unexplained.

Our previous results showed that the Fed's perceptions of inflation's sources relate to policy decisions and yield-curve movements in line with standard macro reasoning. We next ask whether these imbalances merely reflect pre-meeting macro conditions or contain incremental predictive content. Table 7 re-estimates the auxiliary regressions for the policy rate, yields, and term premia with macroeconomic news controls. As before, we also control for the log output gap, core inflation, and the level of the dependent variable at $t-2$. Estimates are qualitatively unchanged, and in most cases stronger. Demand imbalances are more strongly related to the policy rate and to the risk neutral yields. Supply imbalances are more weakly related to the policy rate and risk neutral yields, but more related to term premia.

4.7 Monetary Policy Shocks

Next, we estimate how the imbalances map into asset-price movements around FOMC announcements. Table 8 reports OLS regressions of four-day changes (one day before to

Table 7: Interest rates, yields, and term premia's response controlling for macro news

	FFR/ Shadow Rate	*Risk Neutral Yields*		*Term Premia*	
	i_t^s	1y	5y	1y	5y
I_s	0.082*	0.045	0.024	0.024*	0.046***
	(0.044)	(0.032)	(0.030)	(0.014)	(0.017)
I_d	0.150**	0.203***	0.197***	-0.004	0.030
	(0.074)	(0.065)	(0.064)	(0.018)	(0.020)
Obs.	284	284	284	284	284
R^2	0.983	0.979	0.965	0.882	0.971
Controls	✓	✓	✓	✓	✓

Notes: The table re-estimates regressions of the policy rate, risk-neutral yields, and term premia on demand and supply imbalances, with and without the set of macro-news predictors of Bauer and Swanson (2023). Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and the level of the dependent variable at $t-2$. See the main text for details on the variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

two days after) in risk-neutral yields and term premia on demand and supply imbalances, with and without the macro-news controls. We also estimate analogous regressions of the monetary policy shocks of Romer and Romer (2004), updated by Acosta (2023), on the same imbalances, again with and without these controls.[23]

The earlier asymmetry persists: changes in term premia load positively on supply imbalances, while changes in risk-neutral yields load positively on demand imbalances, although the latter is insignificant at longer maturities. With controls, the one-year risk-neutral yield change also loads negatively on supply at the 10% level. The Romer and Romer monetary policy shock loads positively on both imbalances. With the macro-news controls, both coefficients remain positive, with supply significant at 10% and demand at 5%.[24]

High-frequency monetary policy shocks. A large literature in monetary economics constructs policy surprises from changes in federal funds futures, Eurodollar futures, bond yields, and other asset prices in a narrow window around FOMC announcements.[25] At the

[23] Data from Miguel Acosta's website.

[24] Appendix Tables D.12 and D.16 report results under the alternative normalizations described in Appendix F.1. Results are very similar across these specifications.

[25] Kuttner (2001) uses a scaled version of the current-month federal funds futures (FF1) as his monetary policy shock. Nakamura and Steinsson (2018) and Gürkaynak, Sack, and Swanson (2005) use scaled versions of FF1 and next-month federal funds futures (FF2), along with Eurodollar futures contracts expiring in two to four quarters (ED2, ED3, and ED4). Swanson (2021) includes S&P 500 futures, FF1, FF2, ED1 to ED4, and

Table 8: Imbalances on monetary policy shocks around FOMC announcements

	Δ *Risk-Neutral Yields*				Δ *Term Premia*				*Romer-Romer MPS*	
	1y	1y	5y	5y	1y	1y	5y	5y	MPS	MPS
I_d	0.024**	0.021**	0.015	0.014	-0.008*	-0.006	-0.000	0.002	0.040	0.054**
	(0.009)	(0.010)	(0.010)	(0.012)	(0.005)	(0.005)	(0.007)	(0.009)	(0.029)	(0.025)
I_s	-0.008	-0.013*	-0.005	-0.012	0.012***	0.013***	0.017***	0.019***	0.029	0.031*
	(0.006)	(0.007)	(0.008)	(0.009)	(0.004)	(0.005)	(0.005)	(0.007)	(0.023)	(0.017)
Obs.	284	284	284	284	284	284	284	284	245	245
R^2	0.021	0.078	0.007	0.043	0.029	0.032	0.018	0.039	0.028	0.199
Controls	✗	✓	✗	✓	✗	✓	✗	✓	✗	✓

Notes: The table reports regressions of monetary policy shocks, changes in risk-neutral yields, and changes in term premia on supply and demand imbalances, with and without the macro news predictors of Bauer and Swanson (2023). The monetary policy shock series is from Romer and Romer (2004), updated by Acosta (2023). Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Changes are measured as the difference between one day before and two days after each FOMC meeting. Regressions without controls use the same sample as those with controls to facilitate comparison. See the main text for variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

same time, minutes and transcripts are not released immediately after the FOMC meeting; instead, market participants observe FOMC statements.

To reflect information market participants observe right after the meeting, we apply our methodology to FOMC statements. The procedure closely mirrors that presented in Section 2; Appendix A.4 presents the complete prompt used and more details for the processing of the statements.

After identifying and classifying reasons for inflation in the statements, we construct normalized demand and supply imbalance measures as before. Appendix Table D.6 shows that statement discussions align with the more detailed deliberations in minutes and transcripts: statement demand (supply) imbalances load positively and significantly only on minutes demand (supply) imbalances. The R^2 values suggest two additional facts. First, imbalances from the full deliberations relate to what is communicated to the public, but this mapping is noisy. Second, the regression with statement demand imbalances as the dependent variable has an R^2 roughly twice that for statement supply imbalances, indicating that, conditional on being discussed in the meeting, demand considerations are more likely than supply to be reflected in the statements.

We consider five series: changes in the S&P 500 futures, changes in the current-quarter

Treasury yields with maturities of 2, 5, 10, and 30 years, among other series. Although these studies differ in the exact set of series chosen and the methods used to aggregate them into factors, most of them rely on high-frequency changes in asset prices around FOMC announcements. See Brennan et al. (2024) for a comparison of the data and methods used in this literature.

and three-quarter-ahead Eurodollar future contracts, changes in VIXY, and the monetary policy shock from Bauer and Swanson (2023). All of these changes are taken between ten minutes before and twenty minutes after the FOMC announcement.[26] We regress each of these variables on the demand and supply imbalances constructed using the statements, starting May, 1999.[27] Table 9 presents the results of these regressions. In addition, it also reports the results of the regressions controlling for the set of macro news predictors described previously in Section 4.6.

Table 9: High-Frequency Monetary Policy Shock Regressions

	MPS		*S&P 500*		*ED1*		*ED4*		*VIXY*	
$I^{stat}_{d,t}$	0.009	0.003	-0.048	-0.033	0.006	0.001	0.009	0.002	0.000	-0.002
	(0.006)	(0.004)	(0.047)	(0.048)	(0.004)	(0.003)	(0.008)	(0.006)	(0.002)	(0.002)
$I^{stat}_{s,t}$	0.018***	0.014*	-0.190***	-0.175***	0.015***	0.013**	0.017**	0.011	0.009***	0.009***
	(0.007)	(0.008)	(0.063)	(0.059)	(0.006)	(0.007)	(0.008)	(0.009)	(0.003)	(0.002)
Obs,	178	178	178	178	178	178	178	178	95	95
R^2	0.038	0.183	0.034	0.056	0.042	0.117	0.021	0.180	0.082	0.143
Controls	✗	✓	✗	✓	✗	✓	✗	✓	✗	✓

Notes: The table reports regressions of high-frequency changes in asset prices and the monetary policy shock of Bauer and Swanson (2023) on the demand and supply imbalances extracted from FOMC statements. Asset price changes are from ten minutes before to twenty minutes after the announcement. The dependent variables include the change in the S&P 500 e-mini futures, current-quarter and three-quarter-ahead Eurodollar futures (ED1 and ED4), VIXY, and the monetary policy shock from Bauer and Swanson (2023). VIXY is an exchange-traded fund tracking short-term VIX futures; changes in VIXY are computed from TAQ data. Regressions without controls use the same set of meetings as those with controls. The control set corresponds to the macro news predictors in Bauer and Swanson (2023), described in Section 4.6. Sample runs from May, 1999, to December, 2023. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

The table shows once again that demand and supply perceptions conveyed by the Fed in the statements have very different effects on the asset movements shortly after. Supply imbalances have a positive and mostly significant effect on Eurodollar contracts and in the MPS. In the equity markets, supply imbalances predict a significant increase in VIXY and a decrease in S&P 500 futures. By contrast, demand imbalances have a much smaller and insignificant effect on all of the asset price changes considered.

These results are consistent with our theoretical results on optimal monetary policy when inflation is driven by heterogeneous sources. When inflation reflects supply dis-

[26] The changes in the S&P 500, Eurodollar contracts, and the monetary policy shock are from Bauer and Swanson (2023) and available on the San Franscisco Fed website. We use the S &P 500 e-mini futures contract. VIXY is an ETF of short-term VIX futures, available from 2011 on; VIXY changes are constructed manually from TAQ. Starting in January 2023, ED1 and ED4 are the changes for the SOFR futures contracts.

[27] After May 1999, the Fed began releasing a statement after every meeting, regardless of whether policy changed; see Meade and Acosta (2015) for a discussion of the evolution of statement release practices.

turbances, the policymaker faces a tradeoff: raising rates curbs inflation but deepens the slowdown, while a weaker response supports activity but allows inflation to remain high. Due to this tradeoff, supply pressures create additional risk premia for both equity and bonds because the Fed does not fully stabilize the future path of output and inflation. This is reflected in the rise in VIXY and the fall in S&P 500 futures around the announcement. By contrast, when inflation stems from demand, the future path of output and inflation is stable, and these movements are much smaller.

News articles. The Federal Reserve is not the sole arbiter of economic analysis; financial markets form their own views to infer the likely policy stance taken by the Fed. We ask whether media coverage of upcoming FOMC meetings reveals how the public perceives the Fed's inflation assessment and whether markets respond to the Fed's diagnosis of inflation in the statements over and above information already embedded in pre-meeting news coverage.

We collect news articles that discuss the next FOMC meeting and apply our procedure to extract counts of reasons for higher and lower inflation in the same categories as before. We then construct normalized demand and supply imbalance measures, $I^{news}_{d,t}$ and $I^{news}_{s,t}$. Appendix A.5 details the article collection, filtering, and processing, and reports the prompts used to extract reasons. As a sanity check, Appendix Table D.6 regresses statement-based imbalances on news-based imbalances and shows that statement demand (supply) imbalances load only on news demand (supply) imbalances.

The differences $I^{stat}_{d,t} - I^{news}_{d,t}$ and $I^{stat}_{d,t} - I^{news}_{s,t}$ capture gaps between the Fed's assessment of demand and supply contributions to inflation and the public's diagnosis. Appendix Table D.7 regresses the same asset price changes on $I^{stat}_{d,t} - I^{news}_{d,t}$ and $I^{stat}_{d,t} - I^{news}_{s,t}$, controlling for the macro predictors. Results are qualitatively unchanged: a larger gap between the Fed's and the public's perception of supply's contribution predicts lower market prices and higher volatility around the announcement.[28]

Taken together, these results indicate that media coverage tracks, but does not subsume, the Fed's own inflation diagnostics. News-based demand (supply) imbalances are predictive of demand (supply) imbalances in the statements, yet the residual gap between the two—especially for supply—helps explain high-frequency asset price and volatility responses around FOMC announcements. The statements therefore convey incremental information about the Fed's diagnosis of inflation beyond what is already reflected in pre-meeting news.

[28] Appendix Table D.8 reports an alternative specification that regresses the same high-frequency asset price changes on statement imbalances, controlling for news imbalances and the macro predictors; results are qualitatively unchanged.

5 Model Calibration

In this section, we estimate the parameters we introduced in the model presented in Appendix B, ϕ_s and κ_d. We calibrate the rest of the parameters to standard values from the literature. Substituting the model's expression for expected inflation into the policy rate equation yields:

$$i_t = \phi_s \left[\frac{\alpha\rho_s + \frac{\lambda}{\psi}(1-\rho_s)}{\lambda^2 + \alpha(1-\beta\rho_s)} \right] s_t + \kappa_d \frac{1}{\psi} d_t \tag{20}$$

Next, we define the transformed variables:

$$\tilde{s}_t := \left[\frac{\alpha\rho_s + \frac{\lambda}{\psi}(1-\rho_s)}{\lambda^2 + \alpha(1-\beta\rho_s)} \right] s_t \tag{21}$$

$$\tilde{d}_t := \frac{d_t}{\psi} \tag{22}$$

To express our perceived demand and supply measures in the same units as d_t and s_t, we proceed in two steps. First, we define a smoothed proportion recursively as:

$$P_{x,t} = (1-b_x)P_{x,t-1} + b_x I_{x,t}, \quad x = d, s \tag{23}$$

Then we compute the proxy for the perceived contributions of demand and supply as $\tilde{d}_t = P_{d,t} \times \pi_t$ and $\tilde{s}_t = P_{s,t} \times \pi_t$. We calibrate b_d and b_d separately so that the quarterly persistence of d_t and s_t is equal to 0.9, as in Clarida, Galí, and Gertler (2000). This yields $b_d = 0.13$ and $b_s = 0.26$.[29]

We then regress i_t on $\tilde{d}_t$ and $\tilde{s}_t$ to obtain $\hat{\phi}_s$ and $\hat{\kappa}_d$, respectively. We use the shadow rate from Wu and Xia (2016) whenever it is available to address zero lower bound periods, and the federal funds rate otherwise. We use headline inflation to construct d_t and s_t.

We fix the remaining variables at standard values from the literature: $\beta = 0.99$, $\lambda = 0.3$, $\alpha = 1$, and $\psi = 1$.[30] We also use Newey-West standard errors.[31] Table 10 presents the results. From the table, we see that both parameters have the expected positive sign. This

[29] We consider the average of eight meetings per year, such that two meetings correspond to approximately one quarter. Alternatively, setting $b_d = b_s = 0.2$ yields very similar results.

[30] We set ψ, λ, and β to the values in Clarida, Galí, and Gertler (2000). We choose α to approximate the roughly 1:1 optimal ratio between the inflation and output-gap weights in Debortoli et al. (2018). This choice yields a higher weight than in Woodford (2009), which sets the output-gap weight to 0.048. Because we aim to emphasize the dual mandate tradeoff, a higher output-gap weight is more appropriate.

[31] We do not incorporate the uncertainty associated with estimating the first set of parameters (or any of the calibrated parameters) when we estimate the second set of parameters.

Table 10: Model Estimation

	Estimate	s.e.	t-stat
ϕ_s	0.1028	0.023	4.025
κ_d	2.0892	0.509	4.108

Notes: The table reports OLS estimates of the factor loadings κ_d and ϕ_s from Equation (20). We use the Wu and Xia (2016) shadow rate whenever it is available and the federal funds rate otherwise. s_t and d_t are constructed with core inflation. The calibrated parameters are: $\beta = 0.99$, $\lambda = 0.3$, $\alpha = 1$, $\psi = 1$. Standard errors are Newey–West with 4 lags. See the main text for more details about the estimation and the choice of the calibrated parameters.

result is in line with the reduced-form estimates from Table 2.

Macro block responses. Using the estimated coefficients, we compute the model-implied macroeconomic impulse responses. Figure 2 displays three sets of IRFs: the equilibrium policy rate i_{t+h} (top panel), expected inflation $\mathbb{E}_t \pi_{t+h}$ (middle), and expected output gap $\mathbb{E}_t x_{t+h}$ (bottom). The blue line shows a unit demand shock (d_t); the red line shows a unit supply shock (s_t).

Two facts emerge. Interest rates respond more to demand than to supply pressures, while supply pressures have larger and more persistent effects on inflation than demand. Consistent with the model, expected inflation and the output gap do not respond to demand shocks in equilibrium because monetary policy fully offsets them. supply shocks are not fully neutralized, so they move both expected inflation and the output gap in equilibrium. The resulting deviations revert only slowly. These patterns match the model's predictions and the empirical evidence presented earlier.

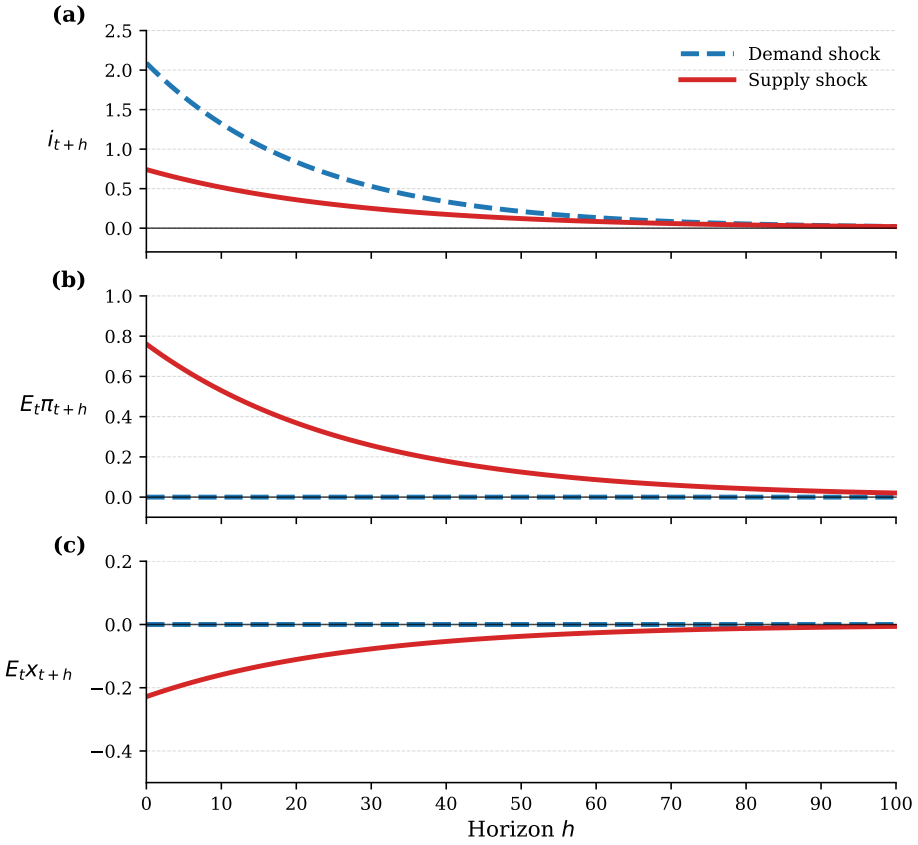
Asset price responses. We now derive the model's asset pricing implications and impulse-response functions. We introduce state-dependent conditional heteroskedasticity for the perceived supply shock innovation, $\varepsilon_{s,t}$, as in Appendix B:

$$Var_t(\varepsilon_{s,t+h}) = \mu_s + \tau_s^h s_t^2. \tag{24}$$

where $\mu_s > 0$ and $\tau_s^h \in (0,1)$. Large supply shocks then raise the conditional variance of future innovations. Appendix B.2 details the derivation and variance calibration, and Figure 3 reports the impulse responses.

We consider a two-period nominal bond and decompose its yield into an expectations-hypothesis term from Equation (11) and an inflation risk-premium term from Equation (12). The top panel shows the response of the expectations-hypothesis term, which is what

Figure 2: Impulse-Response Functions from the Calibrated Model



Notes: The figure plots impulse response functions implied by the model from Section B. Using the estimated parameters $(\hat{\phi}_s, \hat{\kappa}_d)$, calibrated parameters $\beta, \lambda, \alpha, \psi, b_d, b_s$, and the implied persistences ρ_d, ρ_s, we simulate the equilibrium paths of the policy rate i_{t+h} (top panel), expected inflation $\mathbb{E}_t[\pi_{t+h}]$ (middle), and expected output gap $\mathbb{E}_t[x_{t+h}]$ (bottom) following a one-unit innovation in d_t (blue dashed line) or s_t (red solid line). Estimation details and parameter values are provided in the main text. See Appendix B for the details of the model.

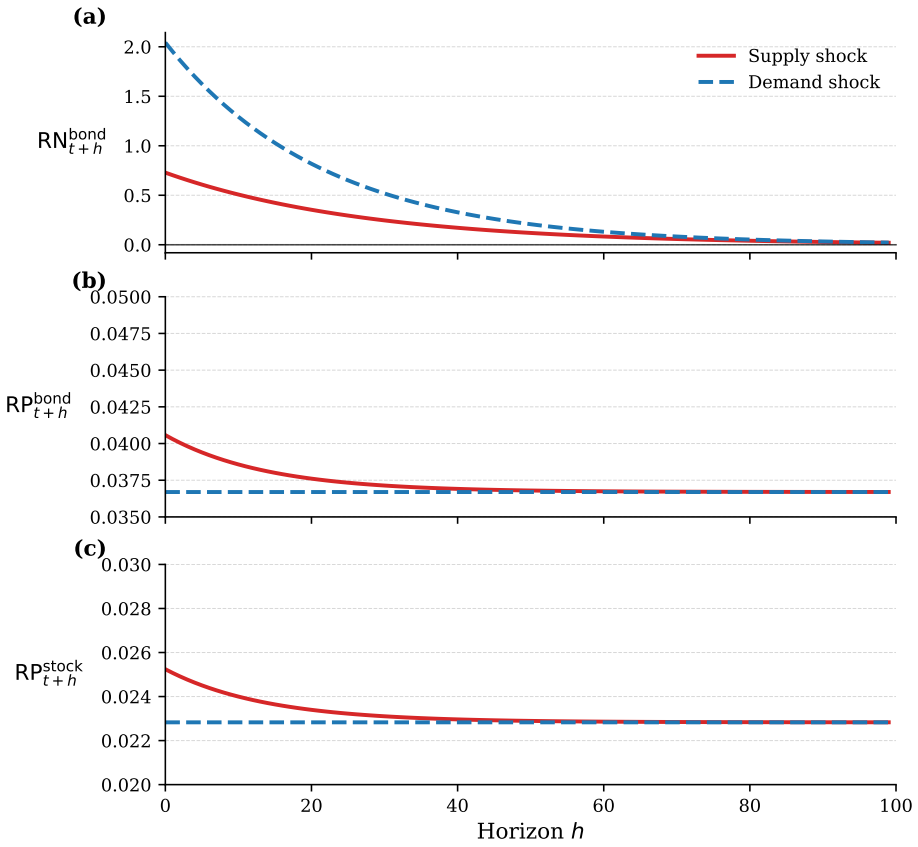
the yield would be if investors were risk-neutral. Because nominal interest rates respond more to demand than to supply pressures, so does the expected path of future rates.

Demand shocks are fully neutralized, so the expected path of nominal rates is unrelated to inflation or consumption. In an economy driven only by demand shocks, bond payoffs are therefore uncorrelated with these variables. Under the optimal policy rule, however, the central bank raises nominal rates in response to supply shocks, but not enough to fully offset inflation. This endogenous response generates comovement between bond prices and inflation, making bonds behave "like stocks," as in Pflueger (2025).

Finally, the bottom panel presents the stock risk premium. In the model, the Fed fully stabilizes the future path of output and inflation after a demand shock, so the stochastic discount factor is uncorrelated with risky-asset returns and demand shocks generate no risk premium. Supply shocks, by contrast, reduce output when consumption is low, so

investors require a positive risk premium.

Figure 3: Impulse Response Functions – Stock and Bond



Notes: The figure plots asset pricing impulse response functions implied by the model from Section B.2. Using the estimated parameters ($\hat{\phi}_s$, $\hat{\kappa}_d$), calibrated parameters $\beta, \lambda, \alpha, \psi, b_d, b_s$, and the implied persistences ρ_d, ρ_s, we simulate the equilibrium paths of the two-period bond yield for a risk-neutral investor (top panel), the bond risk premium (middle), and the stock risk premium (bottom) following a one-unit innovation in d_t (blue dashed line) or s_t (red solid line). In addition, we calibrate a reduced-form state-dependent heteroskedasticity in the perceived supply shock innovation by matching theoretical and empirical moments. Estimation details and parameter values are provided in the main text. See Appendix B for the details of the model.

We conduct an alternative exercise to quantify asset pricing effects. We apply equal-sized impulses to d_t and s_t, scaled so that the impulse to d_t raises interest rates by one percentage point on impact. Under this normalization, d_t increases the expectations-hypothesis term by 2.8 times the increase after s_t. This is similar to the magnitudes found in Table 5 for the effect the demand and supply imbalances on risk neutral yields.

Our stylized model replicates the main empirical patterns. supply pressures raise risk premia on stocks and bonds, whereas demand pressures do not. By contrast, the expected path of future rates responds more to demand than to supply pressures, reflecting the

greater sensitivity of the policy rate to demand in the baseline CGG model.[32]

6 Conclusion

We introduce a novel, LLM–powered narrative methodology that systematically extracts the FOMC's own attributions of inflation and recovers the perceived demand and supply contributions from the full historical record of meeting minutes. The resulting panel of tagged statements opens an information channel that was previously either inaccessible or prohibitively costly to parse at scale. Leveraging this new data set, we document three main results.

Consistent with canonical New Keynesian intuition, realized inflation increases more strongly following supply imbalances than after demand-driven ones. Our historical tag counts therefore recover economically meaningful variation in the shocks that underlie observed price changes.

Augmenting a standard Taylor-rule with separately identified demand and supply factors shows that the funds rate reacts more to demand-driven inflation than to supply-driven inflation. This general pattern is reversed in the pre-Volcker period—policy responds more to supply than to demand—then flips and strengthens in the Volcker–Greenspan era and persists thereafter, consistent with the post-1980 focus on pre-empting demand overheating while partially looking through adverse supply pressures.

We document a systematic asymmetry between the Fed's assessment and yield-curve movements. Higher perceived demand pressure predicts an increase in risk-neutral yields; higher perceived supply pressure predicts an increase in bond and equity premia. This pattern is consistent with New Keynesian logic: demand shocks involve no stabilization trade-off and raise expected future short rates, whereas supply shocks entail a trade-off and elevate interest-rate uncertainty, increasing term premia. These results are robust after controlling for a wide set of variables capturing the state of the economy prior to the FOMC meeting, suggesting that the Fed's assessment carries predictive power beyond merely reflecting past economic data.

[32]The expectations-hypothesis component is the average expectation of future short rates under the physical measure, and abstracts from compensation for bearing duration risk. This differs slightly from the "risk-neutral yield" in Adrian, Crump, and Moench (2013), which is model-implied under the risk-neutral measure. These two objects align if investors were risk-neutral. In our empirical interpretation, revisions to the perceived policy path should therefore primarily appear in the ACM risk-neutral component, while changes in priced duration risk should appear in the ACM term premium. In the model, duration risk arises through covariation of bond payoffs with marginal utility; we focus on the inflation-risk channel because it admits a particularly transparent decomposition (via covariances with expected inflation), even though the empirical term premium may also reflect other sources of risk (e.g., real-rate risk) beyond inflation.

These findings carry several implications. First, the Fed's diagnosis of inflation contains incremental predictive content for both macro outcomes and policy actions; researchers and market participants should therefore track not just the magnitude of inflation but also its perceived origin. Second, policy-rule estimates that ignore the demand–supply distinction risk conflating heterogeneous episodes and understating systematic variation in the reaction function. Third, our LLM-based approach provides a flexible framework for extending narrative identification to other dimensions of central-bank communication, such as financial-stability concerns or distributional objectives, without sacrificing transparency or scalability.

Overall, by bridging modern natural-language processing with textbook monetary economics, we both improve our understanding of inflation and clarify how policymakers adjust to it. We hope the methodology will spur further research into the subtleties contained in central-bank communication and their macro-financial consequences.

References

Acosta, Miguel. 2023. "A New Measure of Central Bank Transparency and Implications for the Effectiveness of Monetary Policy." *International Journal of Central Banking*, 19(3): 49–97.

Adrian, Tobias, Richard K. Crump, and Emanuel Moench. 2013. "Pricing the term structure with linear regressions." *Journal of Financial Economics*, 110(1): 110–138.

Aruoba, S. Borağan and Thomas Drechsel. 2024. "Identifying Monetary Policy Shocks: A Natural Language Approach." National Bureau of Economic Research Working Paper 32417.

Baqaee, David and Emmanuel Farhi. 2022. "Supply and demand in disaggregated Keynesian economies with an application to the Covid-19 crisis." *American Economic Review*, 112(5): 1397–1436.

Basu, Susanto and Brent Bundick. 2017. "Uncertainty Shocks in a Model of Effective Demand." *Econometrica*, 85(3): 937–958.

Bauer, Michael D. and Eric T. Swanson. 2023. "A Reassessment of Monetary Policy Surprises and High-Frequency Identification." *NBER Macroeconomics Annual*, 37: 87–155.

Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam. 2024. "Perceptions about Monetary Policy." *Quarterly Journal of Economics*, 139(4).

Beaudry, Paul, Thomas J Carter, and Amartya Lahiri. 2023. "The Central Bank's Dilemma: Look Through Supply Shocks or Control Inflation Expectations?" National Bureau of Economic Research Working Paper 31741.

Bennani, Hamza, Tobias Kranz, and Matthias Neuenkirch. 2018. "Disagreement between FOMC members and the Fed's staff: New insights based on a counterfactual interest rate." *Journal of Macroeconomics*, 58: 139–153.

Bernanke, Ben. 2022. *21st Century Monetary Policy.* W. W. Norton Company, Incorporated.

Bernanke, Ben S. 2003. "An Unwelcome Fall in Inflation?" *Remarks before the Economics Roundtable, University of California, San Diego, La Jolla, California*.

Black, Fischer. 1995. "Interest Rates as Options." *The Journal of Finance*, 50(5): 1371–1376.

Blinder, Alan S. 2022. *A monetary and fiscal history of the United States, 1961-2021.* Princeton University Press.

Brennan, Connor M., Margaret M. Jacobson, Christian Matthes, and Todd B. Walker. 2024. "Monetary Policy Shocks: Data or Methods?" *Finance and Economics Discussion Series*, , (2024–011): 1–49.

Campbell, John Y., Carolin Pflueger, and Luis M. Viceira. 2020. "Macroeconomic Drivers of Bond and Equity Risks." *Journal of Political Economy*, 128(8): 3148–3185.

Carvalho, Carlos, Fernanda Nechio, and Tiago Tristão. 2021. "Taylor rule estimation by OLS." *Journal of Monetary Economics*, 124: 140–154.

Cieslak, Anna and Carolin Pflueger. 2023. "Inflation and Asset Returns." *Annual Review of Financial Economics*, 15(1): 433–448.

Cieslak, Anna and Michael McMahon. 2023. "Tough Talk: The Fed and the Risk Premium." Working paper.

Cieslak, Anna, Stephen Hansen, Michael McMahon, and Song Xiao. 2023. "Policymakers' uncertainty." National Bureau of Economic Research.

Clarida, Richard, Jordi Galí, and Mark Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature*, 37(4): 1661–1707.

Clarida, Richard, Jordi Galí, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *The Quarterly Journal of Economics*, 115(1): 147–180.

Cogley, Timothy and Thomas J Sargent. 2001. "Evolving post-world war II US inflation dynamics." *NBER macroeconomics annual*, 16: 331–373.

Cogley, Timothy and Thomas J Sargent. 2005. "Drifts and volatilities: monetary policies and outcomes in the post WWII US." *Review of Economic dynamics*, 8(2): 262–302.

Coibion, Olivier and Yuriy Gorodnichenko. 2012. "Why Are Target Interest Rate Changes so Persistent?" *American Economic Journal: Macroeconomics*, 4(4): 126–162.

Curti, Filippo and Sophia Kazinnik. 2023. "Let's face it: Quantifying the impact of nonverbal communication in FOMC press conferences." *Journal of Monetary Economics*, 139: 110–126.

Debortoli, Davide, Jinill Kim, Jesper Lindé, and Ricardo Nunes. 2018. "Designing a Simple Loss Function for Central Banks: Does a Dual Mandate Make Sense?" *The Economic Journal*, 129(621): 2010–2038.

di Giovanni, Julian, Sebnem Kalemli-Ozcan, Alvaro Silva, and Muhammed Yildirim. 2023. "Global Supply Chain Pressures, Trade, and Inflation." Working paper / conference slides.

Eichengreen, Barry and Peter M. Garber. 1991. "Before the Accord: U.S. Monetary-Financial Policy, 1945-51."

Eickmeier, Sandra and Boris Hofmann. 2022. "What drives inflation? Disentangling demand and supply factors." Bank for International Settlements. Forthcoming in *International Journal of Central Banking*.

Engelberg, Joseph, Asaf Manela, William Mullins, and Luka Vulicevic. 2025. "Entity Neutering."

Fernández-Villaverde, Jesús, Pablo Guerrón-Quintana, Juan F Rubio-Ramírez, and Martin Uribe. 2011. "Risk Matters: The Real Effects of Volatility Shocks." *American Economic Review*, 101(6): 2530–2561.

Friedman, Milton and Anna Jacobson Schwartz. 1963. *A Monetary History of the United States, 1867-1960.* Princeton University Press.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. "Text as data." *Journal of Economic Literature*, 57(3): 535–74.

Goodfriend, Marvin. 1991. "Interest rates and the conduct of monetary policy." *Carnegie-Rochester Conference Series on Public Policy*, 34: 7–30.

Goodfriend, Marvin and Robert G. King. 2005. "The incredible Volcker disinflation." *Journal of Monetary Economics*, 52(5): 981–1015.

Gorodnichenko, Yuriy. 2021. "Comment on "The People versus the Markets: A Parsimonious Model of Inflation Expectations"." *Brookings Papers on Economic Activity*, Fall: 370–377.

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera. 2023. "The Voice of Monetary Policy." *American Economic Review*, 113(2): 548–584.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Ivan Werning. 2021. "Monetary Policy in Times of Structural Reallocation." *SSRN Electronic Journal*.

Gürkaynak, Refet S, Brian Sack, and Eric Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking*, 1(1).

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2007. "The U.S. Treasury yield curve: 1961 to the present." *Journal of Monetary Economics*, 54(8): 2291–2304.

Hansen, Anne Lundgaard and Sophia Kazinnik. 2023. "Can ChatGPT Decipher Fedspeak?" *SSRN Electronic Journal*.

Hansen, Stephen, Michael McMahon, and Andrea Prat. 2017. "Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach." *The Quarterly Journal of Economics*, 133(2): 801–870.

Hofmann, Boris, Cristina Manea, and Benoit Mojon. 2024. "Targeted Taylor rules: some evidence and theory." Bank for International Settlements BIS Working Papers 1234.

Kim, Chang-Jin and Charles R Nelson. 2006. "Estimation of a forward-looking monetary policy rule: A time-varying parameter model using ex post data." *Journal of Monetary Economics*, 53(8): 1949–1966.

Kuttner, Kenneth N. 2001. "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market." *Journal of Monetary Economics*, 47(3): 523–544.

Lucca, David and Francesco Trebbi. 2009. *Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements.*

Meade, Ellen E. and Miguel Acosta. 2015. "Hanging on Every Word: Semantic Analysis of the FOMC's Postmeeting Statement." *FEDS Notes*, 2015.0(1580).

Nakamura, Emi and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.

Nakamura, Emi, Venance Riblier, and Jón Steinsson. 2025. "Beyond the Taylor Rule." Working Paper.

Orphanides, Athanasios. 2002. "Monetary-Policy Rules and the Great Inflation." *American Economic Review*, 92(2): 115–120.

Orphanides, Athanasios. 2003. "Historical monetary policy analysis and the Taylor rule." *Journal of Monetary Economics*, 50(5): 983–1022.

Pflueger, Carolin. 2025. "Back to the 1980s or not? The drivers of inflation and real risks in Treasury bonds." *Journal of Financial Economics*, 167: 104027.

Pflueger, Carolin and Gianluca Rinaldi. 2022. "Why does the Fed move markets so much? A model of monetary policy and time-varying risk aversion." *Journal of Financial Economics*, 146(1): 71–89.

Reis, Ricardo. 2022*a*. "The burst of high inflation in 2021â€"22: how and why did we get here?" Bank for International Settlements BIS Working Papers 1060.

Reis, Ricardo. 2022*b*. "Losing the Inflation Anchor." *Brookings Papers on Economic Activity*, 2021(2): 307–379.

Romer, Christina and David Romer. 1989. "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." National Bureau of Economic Research, Inc NBER Working Papers 2966.

Romer, Christina D and David H Romer. 2002. "A Rehabilitation of Monetary Policy in the 1950's." *American Economic Review*, 92(2): 121–127.

Romer, Christina D and David H Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94(4): 1055–1084.

Romer, Christina D and David H Romer. 2008. "The FOMC versus the Staff: Where Can Monetary Policymakers Add Value?" *American Economic Review*, 98(2): 230–235.

Romer, Christina D. and David H. Romer. 2023. "Presidential Address: Does Monetary Policy Matter? The Narrative Approach after 35 Years." *American Economic Review*, 113(6): 1395–1423.

Rotemberg, Julio J. 2013. "Penitence after Accusations of Error: 100 Years of Monetary Policy at the U.S. Federal Reserve." National Bureau of Economic Research. Cambridge, MA. Prepared for the NBER Symposium on "The First Hundred Years of the Federal Reserve," July 10, 2013.

Sack, Brian and Volker Wieland. 2000. "Interest-rate smoothing and optimal monetary policy: a review of recent empirical evidence." *Journal of Economics and Business*, 52(1–2): 205–228.

Shapiro, Adam Hale. 2024. "Decomposing Supply- and Demand-Driven Inflation." *Journal of Money, Credit and Banking*.

Shapiro, Adam Hale and Daniel J. Wilson. 2022. "Taking the Fed at its Word: A New Approach to Estimating Central Bank Objectives using Text Analysis." *The Review of Economic Studies*, 89(5): 2768–2805.

Smets, Frank and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

Svensson, Lars E.O. 1999. "Inflation targeting as a monetary policy rule." *Journal of Monetary Economics*, 43(3): 607–654.

Swanson, Eric T. 2021. "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets." *Journal of Monetary Economics*, 118: 32–53.

Taylor, John B. 1993. "Discretion versus policy rules in practice." *Carnegie-Rochester Conference Series on Public Policy*, 39: 195–214.

Wei, Jason, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models."

Woodford, Michael. 2009. *Interest and Prices.* Princeton:Princeton University Press.

Wu, Jing Cynthia and Fan Dora Xia. 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking*, 48(2–3): 253–291.

Appendix

- Section A contains details of the minutes prompt designs for our LLM methodology.
- Section B contains the details of the model.
- Section C contains supplementary figures.
- Section D contains supplementary tables.
- Section E contains examples of the LLM output.
- Section F contains robustness checks and extensions.

A Minute Prompts

A.1 First Stage Prompt

The following prompt identifies each reason discussed by the Fed in the minutes for the cause of inflation. After the prompt, we attach the text for a single FOMC minute. We repeat this process separately for each minute in the sample. We use the GPT-5 API for this task.[33] Since GPT-5 is a reasoning model, it allows no temperature control.

We use the structured output mode of the API, supplying a JSON schema that defines each key's type and admissible values. This enforces a model response that adheres to the structure that we are imposing. For instance, it enforces that that `reason` need to be a string, `speaker` can only be staff, FOMC, both, or `na`, and so on. We also ask for an importance score varying between 0 and 10 and is used to filter reasons.

The prompt is as follows:

First Stage Prompt: Identifying Reasons

```
You are an economist analyzing the following FOMC minutes. Read the text
carefully and, based solely on its content, answer using minimal language and
the abbreviations below:

* Use "Y" for yes, "N" for no, and "na" if unclear.
* For the speaker, use "FOMC" for FOMC participants, "staff" for staff, "both"
for both FOMC participants and staff, and "na" if unclear (e.g., the narration
does not allow identification). Heuristics:
  - Text under sections titled "Staff Review of the Economic Situation/Financial
  Situation" or phrased as "the staff reported/briefed/noted" → "staff".
  - Text under "Participants" / "Committee" sections or phrased as "participants
  noted/judged/remarked" → "FOMC".
  - If the same passage explicitly attributes to both participants and staff →
  "both".
  - If attribution cannot be inferred from the immediate context → "na".

Answer the following:

Reasons for the Deviation
```

[33] More information about the model and the API are available at https://platform.openai.com/docs/models/gpt-5.

- List ALL reasons that FOMC participants, staff, other speakers, or the meeting narration explicitly cite in the text as structural causes or underlying factors driving inflation or the output gap (positively or negatively). Do not include self-referential answers such as "increasing inflation" or just "inflation concerns" without an accompanying structural cause (e.g., "oil prices").
- Include reasons explicitly linked in the text to (i) inflation, inflation expectations, or prices (e.g., explicit mentions of "inflation," "prices," "inflation expectations," "inflation compensation," "breakevens," "CPI," "PCE," "deflator," or "core inflation"), or (ii) the output gap/slack/aggregate demand vs potential (e.g., "slack," "capacity utilization," "overheating," "demand above/below potential").
- **Hard inclusion rule 1 (Inflation OR Output-Gap trigger):** Include a reason only if the quoted sentence or passage that you provide contains either (a) an explicit reference to inflation/prices/expectations (e.g., "inflation", "prices", "inflation expectations", "inflation compensation", "breakevens", "CPI", "PCE", "deflator", "core inflation"), or (b) an explicit reference to the output gap/slack/aggregate demand relative to potential (e.g., "slack," "overheating," "capacity utilization," "below/above potential," "resource slack").
- **Hard Inclusion rule 2:** Include only reasons explicitly linked in the text to U.S. inflation/prices/expectations or the U.S. output gap/slack/aggregate demand relative to potential. Assume references are U.S. unless the sentence explicitly names a foreign country/region; if it does, exclude the reason unless that same sentence explicitly links it to U.S. inflation/prices or to the U.S. output gap (e.g., import prices, exchange-rate pass-through, demand spillovers affecting U.S. slack).

- When a single sentence or passage contains multiple distinct structural causes (e.g., oil prices AND wage increases), extract and list each as a separate entry. Do not merge multiple causes into one entry. Do the same when a passage cites both price/inflation channels and output-gap/slack channels—**unless** they are part of the same causal chain for a single driver (see "No double-counting" rule below).

```
- **No double-counting for joint inflation–gap statements:** If the same
sentence/passage links an output-gap/slack concept directly to inflation **for
one driver**, record **one** reason only—use the most specific structural driver
as the reason (e.g., "tight labor markets" or "employment is tight"), not a
second separate "output gap" reason. Example: if the text says "inflation is
increasing because employment is tight," include **one** entry such as
{"reason": "tight labor markets", ...} and **do not** add a separate "output
gap/slack" entry. However, if the sentence provides **two distinct drivers**
(e.g., "tight labor markets and higher oil prices are boosting inflation"), list
**both** as separate reasons.

- For each reason, include:
  -> "reason": A brief name (e.g., "wage increases", "energy prices"). Prefer
  the most specific structural driver mentioned (e.g., "tight labor markets"
  instead of generic "output gap") to avoid double-counting.
  -> "importance_score": A score from 0 to 10. Higher scores are more critical
  and receive more attention in the FOMC discussion.
  -> "explanation": An explanation that must include the full sentence or
  passage from the minutes that supports the reason (not just a truncated
  phrase). Wrap the quoted part in single quotes (') or escape double quotes so
  that the JSON remains valid. Provide enough context so the quote can be read
  on its own, and ensure this quoted passage itself contains the explicit
  trigger required above (inflation/prices/expectations OR output
  gap/slack/aggregate demand vs potential).
  -> "speaker": Whether the subject who presented the reason was an FOMC
  participant, the Fed staff, or both. If the speaker is not explicitly
  identified or cannot be inferred from the immediate context, use "na".
- If token/space constraints prevent listing every relevant reason, include a
field "all_reasons" with value "N". Otherwise, use "Y".

Return your answer strictly as valid JSON (do NOT wrap the JSON in any markdown
or code fences) using exactly this structure:

{ "reasons": [ { "reason": "<brief name>", "importance_score": <number from 0 to
10>, "explanation": "<short explanation with full supporting quote>", "speaker":
"<FOMC/staff/both/na>" } ], "all_reasons": "<Y/N>" }
```

```
Do not include any markdown formatting or extra text—only output the JSON.

Example Output:

{ "reasons": [ { "reason": "tight labor markets", "importance_score": 7,
"explanation": "Participants noted that 'with labor markets tight, upward
pressure on prices had intensified over recent months'.", "speaker": "FOMC" }, {
"reason": "energy prices", "importance_score": 5, "explanation": "Staff noted
that 'recent increases in energy prices have begun to pass through to consumer
prices, adding to inflation pressures'.", "speaker": "staff" } ], "all_reasons":
"Y" }

Now, here is the FOMC minutes excerpt to analyze:
```

A.2 Second Stage Prompt

The second stage prompt categorizes each reason identified in step 1 (Appendix A.1). We purposefully limit access of the LLM to the whole document, and, instead, provide it only with the label (the `reason` field from step one) and the context (`explanation` field from step one). We again use the GPT-5 API for this task. Like before, we use the structured output mode of the API, supplying a JSON schema that defines each key's type and admissible values. We attach the reason and the explanation at the end of each prompt, separately for each reason-explanation pair (i.e., each observation generated from the first stage).

This task requires the LLM to use textbook economic reasoning to identify the relationship between the inflation source with inflation and output. As discussed in the introduction and in our theoretical model, demand and supply shocks can be distinguished by their different effects on demand and supply. The fields generated in this task are:

1. `output_effect:` The expected effect of the cause on output.
2. `inflation_effect:` The expected effect of the cause on inflation.
3. `category:` A tag which can be either "demand", "supply", "expectations", or "monetary policy", based on the sign restriction implied by the inflation and output effects and by the context. In some rare cases when both demand and supply are present in the passage, the LLM classifies the reason as "Demand & Supply"; this single reason then counts for demand and supply separately. Ambiguous or tautological reasons are labelled "other".

4. `scratch:` A causal chain linking the shock to inflation and output (if such a link exists). This serves a dual purpose: it helps the LLM reason consistently and allows the econometrician to verify the validity of the reasoning, increasing transparency. This is in line with recent literature on machine learning, which has shown that intermediary reasoning significantly helps LLMs produce complex reasoning (Wei et al., 2022).

Additionally, from previous iterations of our prompt engineering, we identified the LLM often had issues with reasons related to labor markets and trade deficits. Therefore, we added additional examples for these and provide guidelines in other edge cases, as well as contrasting examples (i.e., wage increases caused by demand versus supply).

The prompt classifies reasons into monetary policy when the policy stance, action, or communication is seen as the main driver of the reason in discussion. We restrict the LLM to longer-run inflation expectations (over five years) and to only identify changes relative to prior levels. This excludes, for instance, mere mentions that "expectations remain anchored" as a disinflationary reason.

The prompt is as follows:

Second Stage Prompt: Categorizing Reasons

```
You are an economist classifying a reason cited in the FOMC minutes. Use ONLY
the provided text/context (no outside info).

Goal: assign each reason to exactly ONE label from:
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

──────────── CORE LOGIC ────────────
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP (if the mechanism is policy stance/action/communication) >
  Expectations (if the object is expectations/compensation/breakevens) >
```

```
  Supply&Demand (only if BOTH sides are operative for the same market/period in
  the same sentence/adjacent clause) >
  Demand >
  Supply >
  Other.

──────────── CATEGORIES (with concise examples) ────────────
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
                   supply-chain/bottlenecks/shortages/logistics, capacity
                   constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • “Oil prices jumped” → ↑π,↓y
    • “Tariffs raised input costs; firms passed costs to prices” → ↑π,↓y
    • “Dollar depreciation raised import prices” → ↑π,↓y
    • “Supply bottlenecks eased” → ↓π,↑y
    • “COVID lockdowns in China disrupted inputs” → ↑π,↓y
    • “Strike curtailed auto production” → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment,
  fiscal transfers, sentiment, credit/financial conditions, income/wages,
  overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default
  (income/spending channel),
                     UNLESS the SAME sentence explicitly ties wages to cost
                     pass-through/price-setting (then ⇒ Supply).
  Examples:
    • “Consumer spending rebounded on reopening” → ↑π,↑y
    • “Tight labor market with faster nominal wage growth” → ↑π,↑y
    • “Looser financial conditions boosted spending” → ↑π,↑y
    • “Social distancing depressed services demand” → ↓π,↓y
```

```
Supply&Demand
  Use ONLY when the passage itself states BOTH excess demand AND supply
  constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    • “Demand and supply imbalances pushing up prices” → both sides operative

Expectations (long-run only)
  Definition: Longer-run inflation expectations (horizon ≥ 5y; e.g.,
  “long-run/longer-run/long-term,”
              “5y5y forward,” “5–10y survey measure”) that CHANGED relative to
              prior readings/levels/trend.
  Scope restriction A (no change): Do NOT tag as Expectations if the passage
  merely says
              “anchored,” “stable,” “unchanged,” or “little changed” with no
              explicit movement.
  Scope restriction B (horizon filter): Do NOT tag as Expectations if the
  passage refers to short/near-term
              or year-ahead horizons (e.g., “1-year,” “12-month,” “near-term,”
              “next year,” “two-year”).
              Handle those under Demand/Supply if the mechanism is explicit;
              otherwise ⇒ Other.
  Triggers (LR-only): rose/increased/widened/drifted up;
  fell/declined/narrowed/drifted down; shifted higher/lower;
              became unanchored/re-anchored — when explicitly tied to longer-run
              measures.
  Examples:
    • “5y5y inflation compensation rose” → ↑π (mechanism: long-run expectations)
    • “Survey 5–10y inflation expectations drifted down” → ↓π
    • “Long-run expectations are well anchored” → Other (no change;
    inflation_effect = "na")
    • “One-year-ahead expectations increased” → Not Expectations (apply
    Demand/Supply/Other per mechanism)

MP
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
```

```
    • “Rate hikes will cool demand” → ↓π,↓y
    • “QE eased financial conditions” → ↑π,↑y

Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • “Inflation remained elevated amid broader price pressures” (no driver) →
    na/na
    • “Base effects will lift 12-month inflation temporarily” → na/Up

———————— EDGE-CASE OVERRIDES (clause-level) ————————
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  “passed through to prices”), ⇒ Supply.
Pass-through:
  Treat “pass-through” as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Long-run vs short-run expectations:
  Only horizons ≥ 5y (incl. 5y5y forward, “long-run/longer-run/long-term,” 5–10y
  surveys) qualify for category "Expectations".
  Year-ahead / near-term (≤ 2y) expectations — even if they move — are NOT
  "Expectations"; classify as Demand or Supply if
  the passage specifies a spending or cost-pass-through mechanism; otherwise ⇒
  Other.
Yields vs expectations:
  “Breakevens/inflation compensation/expected inflation” ⇒ Expectations.
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.

———————— OUTPUT ————————
Return strict JSON (no markdown, no extra text):
{
```

```
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.

Constraints:
- Use ONLY the provided text/context.
- Prefer explicit textual cues; if mechanism is unclear, output category
"Other".
- Enumerations are strict: any value outside the allowed set is invalid.

Below is the reason to classify and the surrounding context from the FOMC
minutes.
```

A.3 Expectations

This section describes the prompt and refinements used to construct the panel of reasons linking perceived inflation to expectations, discussed further in Appendix F.3. Appendix A.3.1 implements a restrictive filter that classifies a reason as expectations-based only if it concerns long-run expectations and denotes a change; mere statements that expectations are anchored do not qualify. This is the same prompt used in the baseline analysis and is repeated here for convenience, although expectations imbalances are not discussed in the main text. Appendix A.3.2 implements a less restrictive filter that also treats historical terminology for expectations (such as "inflation psychology") as valid reasons.

A.3.1 More Restrictive Classification focusing on long-term expectation changes

Restrictive Refinement for Expectations

```
You are an economist classifying a reason cited in the FOMC minutes. Use ONLY
the provided text/context (no outside info).

Goal: assign each reason to exactly ONE label from:
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

──────────── CORE LOGIC ────────────
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP (if the mechanism is policy stance/action/communication) >
  Expectations (if the object is expectations/compensation/breakevens) >
  Supply&Demand (only if BOTH sides are operative for the same market/period in
  the same sentence/adjacent clause) >
  Demand >
  Supply >
  Other.

──────────── CATEGORIES (with concise examples) ────────────
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
                    supply-chain/bottlenecks/shortages/logistics, capacity
                    constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • "Oil prices jumped" → ↑π,↓y
```

- “Tariffs raised input costs; firms passed costs to prices” → ↑π,↓y
- “Dollar depreciation raised import prices” → ↑π,↓y
- “Supply bottlenecks eased” → ↓π,↑y
- “COVID lockdowns in China disrupted inputs” → ↑π,↓y
- “Strike curtailed auto production” → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment, fiscal transfers, sentiment, credit/financial conditions, income/wages, overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default (income/spending channel),
                      UNLESS the SAME sentence explicitly ties wages to cost pass-through/price-setting (then ⇒ Supply).
  Examples:
    - “Consumer spending rebounded on reopening” → ↑π,↑y
    - “Tight labor market with faster nominal wage growth” → ↑π,↑y
    - “Looser financial conditions boosted spending” → ↑π,↑y
    - “Social distancing depressed services demand” → ↓π,↓y

Supply&Demand
  Use ONLY when the passage itself states BOTH excess demand AND supply constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    - “Demand and supply imbalances pushing up prices” → both sides operative

Expectations (long-run only)
  Definition: Longer-run inflation expectations (horizon ≥ 5y; e.g., “long-run/longer-run/long-term,”
              “5y5y forward,” “5–10y survey measure”) that CHANGED relative to prior readings/levels/trend.
  Scope restriction A (no change): Do NOT tag as Expectations if the passage merely says
              “anchored,” “stable,” “unchanged,” or “little changed” with no explicit movement.
  Scope restriction B (horizon filter): Do NOT tag as Expectations if the passage refers to short/near-term

```
              or year-ahead horizons (e.g., “1-year,” “12-month,” “near-term,”
              “next year,” “two-year”).
              Handle those under Demand/Supply if the mechanism is explicit;
              otherwise ⇒ Other.
  Triggers (LR-only): rose/increased/widened/drifted up;
  fell/declined/narrowed/drifted down; shifted higher/lower;
              became unanchored/re-anchored — when explicitly tied to longer-run
              measures.
  Examples:
    • “5y5y inflation compensation rose” → ↑π (mechanism: long-run expectations)
    • “Survey 5–10y inflation expectations drifted down” → ↓π
    • “Long-run expectations are well anchored” → Other (no change;
    inflation_effect = "na")
    • “One-year-ahead expectations increased” → Not Expectations (apply
    Demand/Supply/Other per mechanism)

MP
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
    • “Rate hikes will cool demand” → ↓π,↓y
    • “QE eased financial conditions” → ↑π,↑y

Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • “Inflation remained elevated amid broader price pressures” (no driver) →
    na/na
    • “Base effects will lift 12-month inflation temporarily” → na/Up

———————— EDGE-CASE OVERRIDES (clause-level) ————————
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  “passed through to prices”), ⇒ Supply.
Pass-through:
```

```
  Treat “pass-through” as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Long-run vs short-run expectations:
  Only horizons ≥ 5y (incl. 5y5y forward, “long-run/longer-run/long-term,” 5–10y
  surveys) qualify for category "Expectations".
  Year-ahead / near-term (≤ 2y) expectations — even if they move — are NOT
  "Expectations"; classify as Demand or Supply if
  the passage specifies a spending or cost-pass-through mechanism; otherwise ⇒
  Other.
Yields vs expectations:
  “Breakevens/inflation compensation/expected inflation” ⇒ Expectations.
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.

———————— OUTPUT ————————
Return strict JSON (no markdown, no extra text):
{
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.

Constraints:
```

```
- Use ONLY the provided text/context.
- Prefer explicit textual cues; if mechanism is unclear, output category
"Other".
- Enumerations are strict: any value outside the allowed set is invalid.

Below is the reason to classify and the surrounding context from the FOMC
minutes.
```

A.3.2 Less Restrictive Classification Including Inflation Psychology Discussion

Capturing Broader Inflation Expectations Discussion

```
You are an economist classifying a reason cited in the FOMC minutes. Use ONLY
the provided text/context (no outside info).

Goal: assign each reason to exactly ONE label from:
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

──────────── CORE LOGIC ────────────
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP (if the mechanism is policy stance/action/communication) >
  Expectations (if the object is inflation expectations/inflation
  psychology/compensation/breakevens) >
  Supply&Demand (only if BOTH sides are operative for the same market/period in
  the same sentence/adjacent clause) >
  Demand >
  Supply >
  Other.

──────────── CATEGORIES (with concise examples) ────────────
```

```
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
                    supply-chain/bottlenecks/shortages/logistics, capacity
                    constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • “Oil prices jumped” → ↑π,↓y
    • “Tariffs raised input costs; firms passed costs to prices” → ↑π,↓y
    • “Dollar depreciation raised import prices” → ↑π,↓y
    • “Supply bottlenecks eased” → ↓π,↑y
    • “COVID lockdowns in China disrupted inputs” → ↑π,↓y
    • “Strike curtailed auto production” → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment,
  fiscal transfers, sentiment, credit/financial conditions, income/wages,
  overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default
  (income/spending channel),
                      UNLESS the SAME sentence explicitly ties wages to cost
                      pass-through/price-setting (then ⇒ Supply),
                      OR explicitly ties wages to inflation
                      expectations/psychology (then ⇒ Expectations).
  Examples:
    • “Consumer spending rebounded on reopening” → ↑π,↑y
    • “Tight labor market with faster nominal wage growth” → ↑π,↑y
    • “Looser financial conditions boosted spending” → ↑π,↑y
    • “Social distancing depressed services demand” → ↓π,↓y

Supply&Demand
  Use ONLY when the passage itself states BOTH excess demand AND supply
  constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    • “Demand and supply imbalances pushing up prices” → both sides operative
```

```
Expectations
  Definition: Inflation expectations or “inflationary psychology”
  (beliefs/attitudes about future inflation),
              typically longer-run (e.g., “long-run/longer-run/long-term,” “5y5y
              forward,” “5–10y survey measure”)
              or NOT explicitly tied to a short horizon, that CHANGED relative
              to prior readings/levels/trend.
  Scope restriction A (anchored/no-change): If the passage merely says
              “anchored,” “stable,” “unchanged,” “little changed,” “no inflation
              psychology,” or
              “no evidence of inflationary expectations” with no explicit
              movement in inflation or activity,
              still tag as Expectations but set inflation_effect = "na" and
              output_effect = "na"
              (do NOT infer a downward inflation effect from anchoring/absence
              alone).
  Scope restriction B (horizon filter): Do NOT tag as Expectations if the
  passage explicitly refers to short/near-term
              or year-ahead horizons (e.g., “1-year,” “12-month,” “near-term,”
              “next year,” “two-year,” “next few quarters”).
              Handle those under Demand/Supply if the mechanism is explicit;
              otherwise ⇒ Other.
  Triggers: rose/increased/widened/drifted up; fell/declined/narrowed/drifted
  down; shifted higher/lower;
              became unanchored/re-anchored — when explicitly tied to inflation
              expectations or inflationary psychology.
  Examples:
    • “5y5y inflation compensation rose” → ↑π (mechanism: expectations)
    • “Survey 5–10y inflation expectations drifted down” → ↓π
    • “Long-run expectations are well anchored” → Expectations (no change;
    inflation_effect = "na")
    • “One-year-ahead expectations increased” → Not Expectations (apply
    Demand/Supply/Other per mechanism)
    • “Inflationary psychology has strengthened” → ↑π (mechanism: expectations)

MP
```

```
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
    • “Rate hikes will cool demand” → ↓π,↓y
    • “QE eased financial conditions” → ↑π,↑y

Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • “Inflation remained elevated amid broader price pressures” (no driver) →
    na/na
    • “Base effects will lift 12-month inflation temporarily” → na/Up

———————— EDGE-CASE OVERRIDES (clause-level) ————————
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  “passed through to prices”), ⇒ Supply; if SAME sentence links wages to
  inflation expectations/psychology, ⇒ Expectations.
Pass-through:
  Treat “pass-through” as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Long-run vs short-run expectations:
  Treat inflation expectations or inflationary psychology as Expectations by
  default when the horizon is not explicitly short (≤ 2y).
  Year-ahead / near-term (≤ 2y) expectations — when explicitly mentioned (e.g.,
  “1-year,” “12-month,” “next year,” “next few quarters,” “near-term,”
  “short-term”) — are NOT "Expectations"; classify as Demand or Supply if
  the passage specifies a spending or cost-pass-through mechanism; otherwise ⇒
  Other.
Yields vs expectations:
  “Breakevens/inflation compensation/expected inflation” ⇒ Expectations.
```

```
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.

──────────── OUTPUT ────────────
Return strict JSON (no markdown, no extra text):
{
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.

Constraints:
- Use ONLY the provided text/context.
- Prefer explicit textual cues; if mechanism is unclear, output category
"Other".
- Enumerations are strict: any value outside the allowed set is invalid.

Below is the reason to classify and the surrounding context from the FOMC
minutes.
```

A.4 Statement Prompts

This appendix reports the prompts used for the FOMC statements. As before, we use a two-stage procedure with prompts that closely match those for the minutes. We also flag whether each reason is explicitly linked to inflation in the text or whether the link

is inferred by the LLM. Because statements are much shorter, we do not filter reasons by importance and do not restrict the sample to explicit reasons in our baseline specification.

First Stage Prompt: Statements

```
You are an economist analyzing the following FOMC statement. Your task is to
extract only structural reasons that the statement links to
inflation/prices/expectations or to the output gap/slack/aggregate demand as
discussed in the statement. Use minimal language and the abbreviations below.

* Use "Y" or "N" only for the all_reasons field.

Answer the following:

Reasons for the Deviation
- List ALL reasons that the statement explicitly cites as structural causes or
underlying factors driving inflation or the output gap (positively or
negatively) as presented in this statement. Do not include self-referential
answers such as "increasing inflation" or just "inflation concerns" without an
accompanying structural cause (e.g., "oil prices").
- Include reasons explicitly linked in the text to (i) inflation, inflation
expectations, or prices (e.g., explicit mentions of "inflation," "prices,"
"inflation expectations," "inflation compensation," "breakevens," "CPI," "PCE,"
"deflator," or "core inflation"), or (ii) the output gap/slack/aggregate demand
vs potential (e.g., "slack," "capacity utilization," "overheating," "demand
above/below potential").
- **Hard inclusion rule 1 (Inflation OR Output-Gap trigger):** Include a reason
only if the quoted sentence/passage you provide contains either (a) an explicit
reference to inflation/prices/expectations, or (b) an explicit reference to the
output gap/slack/aggregate demand relative to potential.
- **Hard inclusion rule 2 (U.S. scope):** Include only reasons explicitly linked
in the text to U.S. inflation/prices/expectations or the U.S. output
gap/slack/aggregate demand. Assume U.S. unless the sentence names a foreign
country/region; if it does, exclude unless that same sentence explicitly links
it to U.S. inflation/prices or to the U.S. output gap (e.g., import prices,
exchange-rate pass-through, demand spillovers affecting U.S. slack).
```

```
- When a single sentence/passage contains multiple distinct structural causes (e.g., oil prices AND wage increases), list each as a separate entry. Do not merge multiple causes.
- **No double-counting for joint inflation–gap statements:** If the same sentence/passage links an output-gap/slack concept directly to inflation for one driver, record one reason only—use the most specific structural driver (e.g., "tight labor markets") and do not add a separate "output gap" reason. If the sentence provides two distinct drivers (e.g., "tight labor markets and higher oil prices are boosting inflation"), list both.

For each reason, include:
  -> "reason": A brief name (e.g., "wage increases", "energy prices"). Prefer
  the most specific structural driver mentioned (e.g., "tight labor markets"
  instead of "output gap").
  -> "importance_score": A score from 0 to 10 reflecting the statement's
  emphasis on how much this driver matters at this time (not in general).
  -> "explanation": A short explanation that includes the full sentence or
  passage supporting the reason. Wrap the quoted part in single quotes (') or
  escape double quotes so JSON remains valid. Ensure the quoted passage itself
  contains the trigger from Hard inclusion rule 1.

Token constraints:
- If space prevents listing every relevant reason, set "all_reasons" = "N". If you listed all reasons found, set "all_reasons" = "Y".
- If the statement contains no reasons that meet Hard inclusion rules 1–2, return an empty "reasons" array and set "all_reasons" = "Y".

Return your answer strictly as valid JSON (do NOT wrap the JSON in markdown/code fences) using exactly this structure:

{ "reasons": [ { "reason": "<brief name>", "importance_score": <number from 0 to 10>, "explanation": "<short explanation with full supporting quote>" } ], "all_reasons": "<Y/N>" }

Example Output:
```

```
{ "reasons": [ { "reason": "tight labor markets", "importance_score": 7,
"explanation": "The statement notes that 'labor markets remain tight, which is
keeping upward pressure on prices and inflation expectations'." }, { "reason":
"energy prices", "importance_score": 5, "explanation": "It also states that
'recent increases in energy prices are expected to keep headline inflation
elevated in coming months'." } ], "all_reasons": "Y" }

Now, here is the statement to analyze:
```

Second Stage Prompt: Statements

```
You are an economist classifying a reason cited in an FOMC statement. Use ONLY
the provided text/context (no outside info).

Goal: assign each reason to exactly ONE label from:
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

———————— EXPLICITNESS FLAG (tag only; do not filter) ————————
Add a field "explicit" ∈ {Y, N, na}:
- Set explicit="Y" if the quoted sentence/passage in the provided context
contains an explicit reference to inflation, prices, or inflation expectations
(e.g., “inflation”, “prices”, “inflation expectations”, “inflation
compensation”, “breakevens”, “CPI”, “PCE”, “deflator”, “core inflation”). If
multiple adjacent sentences form the passage, any one of them containing such a
term is sufficient.
- Set explicit="N" if the reason is only indirectly linked to inflation via
macro reasoning (e.g., output gap, overheating, slack, policy stance) and the
passage itself lacks any explicit inflation/price/expectations term.
- Set explicit="na" only if the context is too truncated/ambiguous to decide.
This flag does NOT change the category or effects—you must still classify
normally.

———————— CORE LOGIC ————————
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
```

```
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP > Expectations > Supply&Demand > Demand > Supply > Other.
———————— CATEGORIES (with concise examples) ————————
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
                    supply-chain/bottlenecks/shortages/logistics, capacity
                    constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • “Oil prices jumped” → ↑π,↓y
    • “Tariffs raised input costs; firms passed costs to prices” → ↑π,↓y
    • “Dollar depreciation raised import prices” → ↑π,↓y
    • “Supply bottlenecks eased” → ↓π,↑y
    • “COVID lockdowns in China disrupted inputs” → ↑π,↓y
    • “Strike curtailed auto production” → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment,
  fiscal transfers, sentiment, credit/financial conditions, income/wages,
  overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default
  (income/spending channel),
                      UNLESS the SAME sentence explicitly ties wages to cost
                      pass-through/price-setting (then ⇒ Supply).
  Examples:
    • “Consumer spending rebounded on reopening” → ↑π,↑y
    • “Tight labor market with faster nominal wage growth” → ↑π,↑y
    • “Looser financial conditions boosted spending” → ↑π,↑y
    • “Social distancing depressed services demand” → ↓π,↓y

Supply&Demand
```

```
  Use ONLY when the passage itself states BOTH excess demand AND supply
  constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    • “Demand and supply imbalances pushing up prices” → both sides operative

Expectations
  Definition: Inflation expectations/anchoring, inflation compensation,
  breakevens/swaps, or firms’ pricing plans BECAUSE expected inflation changed.
  Examples:
    • “Breakevens rose” → ↑π (mechanism: expectations)
    • “Long-run expectations are well anchored” → ↓π risk

MP
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
    • “Rate hikes will cool demand” → ↓π,↓y
    • “QE eased financial conditions” → ↑π,↑y

Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • “Inflation remained elevated amid broader price pressures” (no driver) →
    na/na
    • “Base effects will lift 12-month inflation temporarily” → na/Up

———————— EDGE-CASE OVERRIDES (clause-level) ————————
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  “passed through to prices”), ⇒ Supply.
Pass-through:
  Treat “pass-through” as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
```

```
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Yields vs expectations:
  “Breakevens/inflation compensation/expected inflation” ⇒ Expectations.
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.

———————— OUTPUT ————————
Return strict JSON (no markdown, no extra text):
{
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>",
  "explicit": "<Y | N | na>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.
- explicit: indicator of whether the reason for inflation is explicitly stated
or indirectly inferred by macroeconomic logic.

Constraints:
- Use ONLY the provided text/context.
- Prefer explicit textual cues; if mechanism is unclear, category="Other".
- Enumerations are strict (case-sensitive).
- The explicit flag is a tag, not a filter; classify the reason regardless of
its value.
Below is the reason to classify and the surrounding context from the FOMC
statement.
```

A.5 News Articles and Prompts

The news dataset comes from NexisUni, a popular news database[34]. The data was downloaded using the NexisUni API. The search query follows Gorodnichenko, Pham, and Talavera (2023), and includes the term "FOMC" and at least one of the following terms: "rate" , "monetary", "federal funds rate", "fed funds rate". The time span considered is between May, 1999 and June, 2023. We restrict source type to include "Newspapers", "Newswires Press Releases", and "Webnews".

We further restrict our search to a set of relevant news sources. Our news sample combines general-audience outlets with specialized financial and banking publications. General outlets include Agence France Presse (English), The Telegraph, The Guardian, The Times (London), The Associated Press and The Associated Press International, USA Today, and The New York Times International Edition, which cover the FOMC for a broad readership, emphasizing macroeconomic conditions, political context, and household implications. Specialized outlets include The Bond Buyer, American Banker, and Associated Press Financial Wire, which target market participants and financial institutions and provide more detailed discussion of interest rates, funding conditions, and balance sheet implications. This mix allows us to capture both the public-facing narrative about FOMC decisions and the more technical coverage in the financial press.

We assign each article to the next FOMC date within 30 days; articles that cannot be matched this way are not considered in our analysis. We also de-duplicate articles assigned to an FOMC meeting date. We end up with 1,874 articles. Figure C.9 shows a time-series of the number of articles per year. Figure C.10 presents a wordcloud constructed from the content of these articles.

For each news article, we proceed in two stages as before. We first identify the reasons the article cites as structural causes or underlying factors driving inflation or the output gap that are relevant for the forthcoming FOMC meeting. As before, we require an importance score, and filter reasons with an importance score of at least 5. Then we classify each identified reason into demand, supply, expectations, monetary policy, or other. Additionally, we ask the LLM to provide a flag or whether that reason is explciitly linked to inflation in the text, or whether that link is implicitly inferred by the LLM (e.g., by inferring that a change in the output gap will also affect inflation). In the main exercises, we filter reasons to be explicitly linked to inflation only. We then compute counts of reasons contributing to an increase and decrease of inflation for each FOMC date. Finally, we create normalized imbalances as described in Section 2. The first and second stage prompts

[34] Available at nexisuni.com.

for the news articles are shown below.

First Stage Prompt: News Articles

```
You are an economist analyzing the following news article mentioning a
forthcoming FOMC meeting. Your task is to extract only forward-looking
structural reasons that the article links to inflation/prices/expectations or to
the output gap/slack/aggregate demand for the upcoming FOMC decision. Exclude
retrospective summaries unless the text explicitly ties them to the next
meeting. Use minimal language and the abbreviations below.

* Use "Y" or "N" only for the all_reasons field.

Answer the following:

Reasons for the Deviation
- List ALL reasons that the article explicitly cites as structural causes or
underlying factors driving inflation or the output gap (positively or
negatively) that are relevant for the forthcoming/next FOMC meeting. Do not
include self-referential answers such as "increasing inflation" or just
"inflation concerns" without an accompanying structural cause (e.g., "oil
prices").
- Include reasons explicitly linked in the text to (i) inflation, inflation
expectations, or prices (e.g., explicit mentions of "inflation," "prices,"
"inflation expectations," "inflation compensation," "breakevens," "CPI," "PCE,"
"deflator," or "core inflation"), or (ii) the output gap/slack/aggregate demand
vs potential (e.g., "slack," "capacity utilization," "overheating," "demand
above/below potential").
- **Hard inclusion rule 1 (Inflation OR Output-Gap trigger):** Include a reason
only if the quoted sentence/passage you provide contains either (a) an explicit
reference to inflation/prices/expectations, or (b) an explicit reference to the
output gap/slack/aggregate demand relative to potential.
- **Hard inclusion rule 2 (U.S. scope):** Include only reasons explicitly linked
in the text to U.S. inflation/prices/expectations or the U.S. output
gap/slack/aggregate demand. Assume U.S. unless the sentence names a foreign
country/region; if it does, exclude unless that same sentence explicitly links
it to U.S. inflation/prices or to the U.S. output gap (e.g., import prices,
exchange-rate pass-through, demand spillovers affecting U.S. slack).
```

```
- **Hard inclusion rule 3 (Forward-looking filter):** The quoted
sentence/passage must itself contain a clear forward-looking anchor about the
upcoming/next FOMC decision, such as any of the following cues:
  "ahead of", "before", "going into", "in the run-up to", "at/for the [Month]
  meeting", "next meeting", "upcoming meeting", an explicit date of the next
  meeting, or modal/expectational language tied to the meeting (e.g., "will",
  "likely to", "expected to", "set to", "poised to", "markets are pricing",
  "odds of", "forecast to"). If the structural driver and the forward-looking
  anchor appear in adjacent sentences, include both sentences in the quoted
  passage so that the passage itself contains the forward anchor and the
  trigger.
- **Retrospective exclusion:** Exclude statements that only recap prior meetings
(e.g., "at the last meeting the Fed said…", "the Committee previously noted…")
unless the same sentence explicitly states implications for the next meeting
(e.g., "…which sets the stage for/raises the odds of X at the next meeting").
- When a single sentence/passage contains multiple distinct structural causes
(e.g., oil prices AND wage increases), list each as a separate entry. Do not
merge multiple causes. Do the same when a passage cites both price/inflation
channels and output-gap/slack channels—**unless** they are part of the same
causal chain for a single driver (see "No double-counting" below).
- **No double-counting for joint inflation–gap statements:** If the same
sentence/passage links an output-gap/slack concept directly to inflation for one
driver, record one reason only—use the most specific structural driver (e.g.,
"tight labor markets") and do not add a separate "output gap" reason. If the
sentence provides two distinct drivers (e.g., "tight labor markets and higher
oil prices are boosting inflation"), list both.

For each reason, include:
  -> "reason": A brief name (e.g., "wage increases", "energy prices"). Prefer
  the most specific structural driver mentioned (e.g., "tight labor markets"
  instead of "output gap").
  -> "importance_score": A score from 0 to 10 reflecting the article's emphasis
  on how much this driver matters for the upcoming decision (not in general).
```

```
  -> "explanation": A short explanation that includes the full sentence or
  passage supporting the reason. Wrap the quoted part in single quotes (') or
  escape double quotes so JSON remains valid. Ensure the quoted passage itself
  contains (i) the forward-looking anchor about the next meeting and (ii) the
  trigger from Hard inclusion rule 1. If needed, include two adjacent sentences
  to satisfy both.

Token constraints:
- If space prevents listing every relevant forward-looking reason, set
"all_reasons" = "N". If you listed all forward-looking reasons found, set
"all_reasons" = "Y".
- If the article contains no forward-looking reasons that meet Hard inclusion
rules 1–3, return an empty "reasons" array and set "all_reasons" = "Y".

Return your answer strictly as valid JSON (do NOT wrap the JSON in markdown/code
fences) using exactly this structure:

{ "reasons": [ { "reason": "<brief name>", "importance_score": <number from 0 to
10>, "explanation": "<short explanation with full supporting quote>" } ],
"all_reasons": "<Y/N>" }

Example Output:

{ "reasons": [ { "reason": "tight labor markets", "importance_score": 7,
"explanation": "The article says 'with labor markets still tight ahead of the
September meeting, several analysts expect upward pressure on prices to persist
into Q4,' raising the odds of a hike." }, { "reason": "energy prices",
"importance_score": 5, "explanation": "It adds that 'recent increases in energy
prices are expected to keep headline CPI elevated going into the upcoming
meeting,' which officials will weigh." } ], "all_reasons": "Y" }

Now, here is the article to analyze:
```

Second Stage Prompt: News Articles

```
You are an economist classifying a reason cited in a news article about the FOMC
(often in the run-up to a meeting). Use ONLY the provided text/context (no
outside info).

Goal: assign each reason to exactly ONE label from:
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

──────────── HARD INCLUSION (news) ────────────
- **Hard inclusion rule 1 (Explicit inflation reference required):** A reason
may be included only if the quoted sentence or passage that you provide contains
at least one explicit reference to inflation, prices, or inflation expectations
(e.g., "inflation", "prices", "inflation expectations", "inflation
compensation", "breakevens", "CPI", "PCE", "deflator", "core inflation"). If
this explicit reference is not present in the quoted sentence/passage itself,
classify as "Other" with output_effect="na" and inflation_effect="na". Do NOT
rely on macro inferences (e.g., Phillips curve) to bridge the link.

──────────── CORE LOGIC ────────────
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP (if the mechanism is policy stance/action/communication) >
  Expectations (if the object is expectations/compensation/breakevens) >
  Supply&Demand (only if BOTH sides are operative for the same market/period in
  the same sentence/adjacent clause) >
  Demand >
  Supply >
  Other.

──────────── CATEGORIES (with concise examples) ────────────
```

```
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
                    supply-chain/bottlenecks/shortages/logistics, capacity
                    constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • “Oil prices jumped” → ↑π,↓y
    • “Tariffs raised input costs; firms passed costs to prices” → ↑π,↓y
    • “Dollar depreciation raised import prices” → ↑π,↓y
    • “Supply bottlenecks eased” → ↓π,↑y
    • “COVID lockdowns in China disrupted inputs” → ↑π,↓y
    • “Strike curtailed auto production” → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment,
  fiscal transfers, sentiment, credit/financial conditions, income/wages,
  overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default
  (income/spending channel),
                      UNLESS the SAME sentence explicitly ties wages to cost
                      pass-through/price-setting (then ⇒ Supply).
  Examples:
    • “Consumer spending rebounded on reopening” → ↑π,↑y
    • “Tight labor market with faster nominal wage growth” → ↑π,↑y
    • “Looser financial conditions boosted spending” → ↑π,↑y
    • “Social distancing depressed services demand” → ↓π,↓y

Supply&Demand
  Use ONLY when the passage itself states BOTH excess demand AND supply
  constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    • “Demand and supply imbalances pushing up prices” → both sides operative

Expectations
```

```
  Definition: Inflation expectations/anchoring, inflation compensation,
  breakevens/swaps, or firms' pricing plans BECAUSE expected inflation changed.
  Examples:
    • "Breakevens rose" → ↑π (mechanism: expectations)
    • "Long-run expectations are well anchored" → ↓π risk

MP
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
    • "Rate hikes will cool demand" → ↓π,↓y
    • "QE eased financial conditions" → ↑π,↑y

Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • "Inflation remained elevated amid broader price pressures" (no driver) →
    na/na
    • "Base effects will lift 12-month inflation temporarily" → na/Up

──────────── EDGE-CASE OVERRIDES (clause-level) ────────────
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  "passed through to prices"), ⇒ Supply.
Pass-through:
  Treat "pass-through" as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Yields vs expectations:
  "Breakevens/inflation compensation/expected inflation" ⇒ Expectations.
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.
```

```
———————— OUTPUT ————————
Return strict JSON (no markdown, no extra text):
{
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.

Constraints:
- Use ONLY the provided text/context.
- Prefer explicit textual cues; if mechanism is unclear, output category
"Other".
- Enumerations are strict: any value outside the allowed set is invalid.
- Apply Hard inclusion rule 1 strictly; if not satisfied, set category="Other"
and both effects="na".

Below is the reason to classify and the surrounding context from the news
article.
```

A.6 Entity Neutering

This section presents the prompts used in the entity neutering exercise adapted from Engelberg et al. (2025) and discussed in Appendix F.4. There are three sets of prompts: masking, paraphrasing and guessing. Masking and paraphrasing have separate prompts for retries. The following prompts are the first masking prompt, and "retry" masking prompt, and first paraphrasing prompt, the "retry" paraphrasing prompt, and the guess-

ing prompts, respectively.

Masking Prompt: First Try

```
Your role is to ANONYMIZE all text that is provided by the user.

You will receive JSON with two fields: reason and explanation.
You must anonymize both fields.

You must anonymize ANY information that could be used to identify (i) the
calendar time period
of the underlying meeting (year, month, quarter, decade, or any famous
historical episode),
or (ii) the identity of any individuals referenced.

You should also ANONYMIZE any other information which one could use to identify
the source
or to make an educated guess about when it occurred.

You MUST anonymize all identifiers, including names of individuals,
roles/titles, locations,
organizations, committee/body names, agencies, programs, policy tools, event
names, and ALL times,
years, dates and ALL numbers and percentages in the text, including units.

These should be replaced with person_x, role_x, location_x, org_x, committee_x,
agency_x,
program_x, policy_tool_x, event_x, time_x, year_x, date_x and number a, b, c,
respectively.
Also replace any website or internet links with link_x.

You should never just delete an identifier; instead, always replace it with an
anonymized analog.

CRITICAL: Preserve the economic meaning and structure needed for classification.

You MUST KEEP:
- whether the mechanism is supply vs demand vs expectations/persistence vs
monetary policy
```

```
- whether it refers to inflation vs real activity and the direction of effects
(up/down/ambiguous)
- generic horizons like "near-term", "medium-term", "longer-run", "past year",
"recent months"
- generic sources like "consumer surveys", "business surveys", "market-based
measures"
- generic actors like "consumers", "households", "firms", "participants",
"staff"

You MUST REMOVE/REPLACE ONLY identifiers:
- exact calendar dates/years/months/quarters
- names of people and named roles tied to specific individuals
- named events/episodes that pin down the time period (e.g., pandemics, wars,
crises) by name
- uniquely identifying survey/institution names (you may keep "consumer survey"
generically)
- precise numeric values (levels, percentages, amounts) unless strictly
necessary for the mechanism

Goal: someone should still be able to classify the passage correctly, but should
not be able to pin down the exact meeting date or specific named
individuals/events.

After you read and ANONYMIZE the text, you should output the anonymized text and
nothing else.

Return STRICT JSON ONLY with exactly these keys:
- reason
- explanation
No extra keys.
```

Masking Prompt: Retry

```
Your role is to ANONYMIZE all text that is provided by the user. You have
already tried to anonymize the following text and failed, leaving details which
were used to identify the text by an expert. Be more aggressive in your efforts
to anonymize and pay more attention, so that you do not leave even the most
minute details in the anonymized text.

You will receive JSON with two fields: reason and explanation.
You must anonymize both fields.

You must anonymize ANY information that could be used to identify (i) the
calendar time period
of the underlying meeting (year, month, quarter, decade, or any famous
historical episode),
or (ii) the identity of any individuals referenced.

You should also ANONYMIZE any other information which one could use to identify
the source
or to make an educated guess about when it occurred.

You MUST anonymize all identifiers, including names of individuals,
roles/titles, locations,
organizations, committee/body names, agencies, programs, policy tools, event
names, and ALL times,
years, dates and ALL numbers and percentages in the text, including units.

These should be replaced with person_x, role_x, location_x, org_x, committee_x,
agency_x,
program_x, policy_tool_x, event_x, time_x, year_x, date_x and number a, b, c,
respectively.
Also replace any website or internet links with link_x.

You should never just delete an identifier; instead, always replace it with an
anonymized analog.

CRITICAL: Preserve the economic meaning and structure needed for classification.
```

```
You MUST KEEP:
- whether the mechanism is supply vs demand vs expectations/persistence vs
monetary policy
- whether it refers to inflation vs real activity and the direction of effects
(up/down/ambiguous)
- generic horizons like "near-term", "medium-term", "longer-run", "past year",
"recent months"
- generic sources like "consumer surveys", "business surveys", "market-based
measures"
- generic actors like "consumers", "households", "firms", "participants",
"staff"

You MUST REMOVE/REPLACE ONLY identifiers:
- exact calendar dates/years/months/quarters
- names of people and named roles tied to specific individuals
- named events/episodes that pin down the time period (e.g., pandemics, wars,
crises) by name
- uniquely identifying survey/institution names (you may keep "consumer survey"
generically)
- precise numeric values (levels, percentages, amounts) unless strictly
necessary for the mechanism

Goal: someone should still be able to classify the passage correctly, but should
not be able to pin down the exact meeting date or specific named
individuals/events.

You have already failed me before, take no chances this time and be aggressive
in censoring the text. After you read and ANONYMIZE the text, you should output
the anonymized text and nothing else.

Return STRICT JSON ONLY with exactly these keys:
- reason
- explanation
No extra keys.
```

Paraphrasing Prompt: First Try

```
You are to paraphrase text ensuring that its core information content is
unchanged, such that
someone who has memorized the text could not recognize it. Make sure to replace
all unique or
identifying words and phrases. Also change the sentence structure so that the
resulting text is
completely different from the original. This includes changing quotes and other
highly recognizable
structures in the text.

You will receive JSON with two fields: reason and explanation.
You must paraphrase+anonymize both fields.

Any information which could ever be used to identify when the meeting occurred
or which individuals
are referenced MUST be anonymized. Absolutely no mention of key words related to
famous historical
events, named regimes, or other specific time anchors can be left in the text.

If you observe any names, proper nouns, aliases, acronyms, dates, agencies,
programs, policy tools,
or recognizable quotes, replace them with anonymous strings.

Anonymity is the prime objective, leave zero hints for someone who has memorized
the text!
Ensure that the paraphrased text has approximately the same number of words as
the original text.

CRITICAL: Preserve the economic meaning and structure needed for classification.

You MUST KEEP:
- whether the mechanism is supply vs demand vs expectations/persistence vs
monetary policy
- whether it refers to inflation vs real activity and the direction of effects
(up/down/ambiguous)
```

```
- generic horizons like "near-term", "medium-term", "longer-run", "past year",
"recent months"
- generic sources like "consumer surveys", "business surveys", "market-based
measures"
- generic actors like "consumers", "households", "firms", "participants",
"staff"

You MUST REMOVE/REPLACE ONLY identifiers:
- exact calendar dates/years/months/quarters
- names of people and named roles tied to specific individuals
- named events/episodes that pin down the time period (e.g., pandemics, wars,
crises) by name
- uniquely identifying survey/institution names (you may keep "consumer survey"
generically)
- precise numeric values (levels, percentages, amounts) unless strictly
necessary for the mechanism

Goal: someone should still be able to classify the passage correctly, but should
not be able to pin down the exact meeting date or specific named
individuals/events.

After you read and ANONYMIZE the text, you should output the anonymized text and
nothing else.

Return STRICT JSON ONLY with exactly these keys:
- reason
- explanation
No extra keys.
```

Paraphrasing Prompt: Retry

```
You are to paraphrase text ensuring that its core information content is
unchanged, such that
someone who has memorized the text could not recognize it. You have already
failed to paraphrase
the text to make it anonymous. Be more aggressive in your paraphrasing and
anonymizing of the text.
```

```
Make sure to replace all unique or identifying words and phrases. Also change
the sentence structure
so that the resulting text is completely different from the original. This
includes changing quotes
and other highly recognizable structures in the text.

You will receive JSON with two fields: reason and explanation.
You must paraphrase+anonymize both fields.

Any information which could ever be used to identify when the meeting occurred
or which individuals
are referenced MUST be anonymized. Absolutely no mention of key words related to
famous historical
events, named regimes, or other specific time anchors can be left in the text.

If you observe any names, proper nouns, aliases, acronyms, dates, agencies,
programs, policy tools,
or recognizable quotes, replace them with anonymous strings.

Anonymity is the prime objective, leave zero hints for someone who has memorized
the text!
Ensure that the paraphrased text has approximately the same number of words as
the original text.

CRITICAL: Preserve the economic meaning and structure needed for classification.

You MUST KEEP:
- whether the mechanism is supply vs demand vs expectations/persistence vs
monetary policy
- whether it refers to inflation vs real activity and the direction of effects
(up/down/ambiguous)
- generic horizons like "near-term", "medium-term", "longer-run", "past year",
"recent months"
- generic sources like "consumer surveys", "business surveys", "market-based
measures"
- generic actors like "consumers", "households", "firms", "participants",
"staff"
```

```
You MUST REMOVE/REPLACE ONLY identifiers:
- exact calendar dates/years/months/quarters
- names of people and named roles tied to specific individuals
- named events/episodes that pin down the time period (e.g., pandemics, wars,
crises) by name
- uniquely identifying survey/institution names (you may keep "consumer survey"
generically)
- precise numeric values (levels, percentages, amounts) unless strictly
necessary for the mechanism

Goal: someone should still be able to classify the passage correctly, but should
not be able to pin down the exact meeting date or specific named
individuals/events.

You have already failed once before, take no chances this time, be aggressive in
your effort to anonymize the text. After you read and ANONYMIZE the text, you
should output the anonymized text and nothing else.

Return STRICT JSON ONLY with exactly these keys:
- reason
- explanation
No extra keys.
```

Guessing Prompt

```
You will receive a body of text which has been anonymized. You are omniscient.
Use all of your knowledge and any clues in the text to infer the most likely
exact meeting date.

IMPORTANT:
- You MUST output a date in YYYY-MM-DD format even if you are unsure.
- If the text provides insufficient evidence, make your best guess and set a low
confidence.
- Do NOT refuse and do NOT output null.

Return STRICT JSON ONLY that matches the schema. Do not include any extra keys.
```

B Model

This section develops a simple structural framework that formalizes how demand and supply disturbances propagate into inflation and trigger monetary-policy responses. This motivates the estimation a dynamic Taylor rule that explicitly conditions these responses on the monetary authority's assessment of the underlying drivers of inflation.

We closely follow the Clarida, Galí, and Gertler (1999) framework (CGG). Let y_t and z_t denote the components for output and the natural level of output (deviation from long-run trend, in logs). We define the output gap as $x_t := y_t - z_t$. Finally, let π_t denote the inflation rate and i_t the nominal interest rate (in deviation from the long-run level).

The model yields two standard behavioral equations, derived from household and firm optimization. The IS equation (IS) follows directly from the household's Euler equation for consumption, while the Phillips curve (PC) arise from firms' optimal price-setting under nominal rigidities (e.g., Calvo pricing):

$$x_t = -\psi(i_t - \mathbb{E}_t\pi_{t+1}) + \mathbb{E}_t x_{t+1} + \kappa_d d_t \tag{IS}$$

$$\pi_t = \lambda x_t + \beta\mathbb{E}_t\pi_{t+1} + \phi_s s_t \tag{PC}$$

Variables d_t and s_t denote the demand and supply disturbances. The positive loadings $\phi_s, \kappa_d > 0$ scale the differential effects of these factors on macro outcomes: ϕ_s maps supply shocks into the output gap, and κ_d maps demand shocks into inflation. Their data generating process is given by:

$$d_t = \rho_d d_{t-1} + \varepsilon_{d,t} \tag{25}$$

$$s_t = \rho_s s_{t-1} + \varepsilon_{s,t} \tag{26}$$

with $\varepsilon_{d,t}$, $\varepsilon_{s,t}$ identically distributed, independent, zero-mean innovations and ρ_d, $\rho_s \in (0,1)$.

According to the equations above, a positive demand shock (demand curve shifted to the right) increases the output gap and inflation directly. It also increases inflation indirectly via the effect of x_t on π_t in Equation (PC). A positive supply shock (supply curve shifted to the left) increases inflation and decreases output; however, it keeps the output gap unchanged because potential output moves one-to-one with the changes in output. Thus, demand and supply shocks are distinguished by their differential impacts on output and inflation. The policy instrument is the short nominal rate (i_t), and all forecasts are common knowledge—there is no disagreement between the monetary authority and market participants.

As shown in the detailed derivations in Appendix B.1, the equilibrium interest rate under the baseline discretionary case will be:

$$i_t = \mathbb{E}_t \pi_{t+1} + \Upsilon_s s_t + \Upsilon_d d_t \tag{27}$$

where Υ_s, Υ_d are functions of the structural parameters only. Factors d_t, s_t are persistent over time. Consequently, decomposing the shock into d_t and s_t informs the central bank's forecasts of future inflation and unemployment and, in turn, the interest-rate path it sets. Moreover, this decomposition enters forward expectations in a rational-expectations equilibrium.

The linear-quadratic policy rule implies a stark asymmetry between demand and supply disturbances. Because the real rate exactly offsets the demand shock $\kappa_d d_t$, the central bank neutralizes shocks to the output gap one-to-one. In addition, the optimal response coefficient on d_t, Υ_d, is always strictly positive. By contrast, the supply shock, $\phi_s s_t$, is only partially offset. As we show in Appendix B.1, under standard parameter choices, we have:

$$0 < \Upsilon_s < \Upsilon_d \tag{28}$$

and real interest rates will respond more to demand shocks than to supply shocks. Hence, monetary policy systematically leans harder against demand-driven inflation than against supply-driven inflation—a prediction that guides our empirical tests below.

B.1 Additional model equilibrium derivations

In this section, we follow Clarida, Galí, and Gertler (1999) (CGG) and derive the optimal policy under rational expectations. The Fed's objective function is to minimize a quadratic loss respecting its dual mandate:

$$\min \frac{1}{2}\mathbb{E}_t\left(\sum_{j=0}^{\infty}\beta^j(\alpha x_{t+j}^2 + \pi_{t+i}^2)\right) \tag{$\star$}$$

We consider the baseline case of discretionary policy without commitment, in which the Fed chooses i_t by reoptimizing every period. That is, the Fed chooses a triplet $\{x_t, \pi_t, i_t\}$ every period t to optimize Equation ($\star$) subject to the supply curve (PC) and the demand curve (IS). Following CGG, the problem can be further simplified by considering a two-step procedure:

1. First, the Fed chooses x_t and π_t to minimize Equation ($\star$) given Equation (PC).
2. Then, given x_t and π_t, it pins down i_t from Equation (IS). This will be the interest rate that supports the pair chosen.

Because of discretion, the Fed cannot manipulate beliefs. Thus, the Fed takes the private sector beliefs as given. The Lagrangian of the first optimization problem will be:

$$\mathcal{L} = -\frac{1}{2}\left[\alpha x_t^2 + \pi_t^2\right] + \gamma(\pi_t - \lambda x_t - f_t) + F_t \tag{29}$$

where

$$f_t = \beta \mathbb{E}_t \pi_{t+1} + \phi_s s_t \tag{30}$$

$$F_t = -\frac{1}{2}\mathbb{E}_t\left[\sum_{j=1}^{\infty} \beta^j (\alpha x_{t+j}^2 + \pi_{t+j}^2)\right] \tag{31}$$

are both treated as fixed in the optimization. The FOC yields:

$$x_t = -\frac{\lambda}{\alpha}\pi_t \tag{32}$$

The next step is to plug this into the the Phillips Curve (PC):

$$\pi_t = -\frac{\lambda^2}{\alpha}\pi_t + \beta \mathbb{E}_t \pi_{t+1} + \phi_s s_t \tag{33}$$

$$\implies \pi_t = \frac{\alpha}{\alpha + \lambda^2}\beta \mathbb{E}_t \pi_{t+1} + \frac{\alpha}{\alpha + \lambda^2}\phi_s s_t \tag{34}$$

To close the model, we need to make some assumption about how the market forms beliefs about the future. We assume that the market shares the same beliefs as the Fed with respect to d_t, s_t and all other variables:

$$\pi_t = \frac{\alpha}{\alpha + \lambda^2}\beta \mathbb{E}_t \pi_{t+1} + \frac{\alpha}{\alpha + \lambda^2}\phi_s s_t \tag{35}$$

Imposing rational expectations and the usual transversality conditions (which rule out

bubbles), we can iterate forward, conditional on the current beliefs about the shocks:

$$\begin{aligned}\pi_t =& \frac{\alpha}{\alpha+\lambda^2}(\phi_s s_t) + \left(\frac{\alpha}{\alpha+\lambda^2}\right)^2 \beta(\phi_s \rho_s s_t) \\ &+ \left[\frac{\alpha}{\alpha+\lambda^2}\right]^3 \beta^2(\phi_s \rho_s^2 s_t) + \ldots \\ =& \frac{\alpha}{\alpha+\lambda^2}\phi_s \left[s_t + \frac{\beta\alpha\rho_s}{\alpha+\lambda^2} s_t + \left(\frac{\beta\alpha\rho_s}{\alpha+\lambda^2}\right)^2 s_t + \ldots\right] \\ =& \frac{\alpha\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)} s_t \end{aligned} \tag{36}$$

The expression above decomposes equilibrium inflation in terms of s_t. From this expression, it also follows that:

$$\mathbb{E}_t \pi_{t+h} = \frac{\alpha\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)} \rho_s^h s_t, \quad \forall h \geq 0 \tag{37}$$

$$x_t = -\frac{\lambda\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)} s_t \tag{38}$$

$$\mathbb{E}_t x_{t+h} = -\frac{\lambda\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)} \rho_s^h s_t, \quad \forall h \geq 0 \tag{39}$$

where the equations above follow from the FOC and the data generating process for d_t and s_t.

Note that the total IS shock ($\kappa_d d_t$) does not appear in the expression for equilibrium inflation. The reason is that the Fed sets interest rates to exactly offset it. To see this, note that by plugging in the FOC solution we can rewrite the IS curve (IS) as:

$$i_t = \mathbb{E}_t \pi_{t+1} + \underbrace{\frac{1}{\psi}\left[\frac{\lambda(1-\rho_s)\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)}\right]}_{\Upsilon_s} s_t + \underbrace{\frac{1}{\psi}\kappa_d}_{\Upsilon_d} d_t \tag{40}$$

The equation above implies that the real interest rate ($i_t - \mathbb{E}_t \pi_{t+1}$) will fully offset the demand shock $\kappa_d d_t$. This can be seen by the last term in i_t, which depends on $\kappa_d d_t/\psi$. Recall that ψ measures the effectiveness of real interest rates (how much a unit interest rate increase lowers the output gap), which means that interest rates will adjust as to entirely offset $\kappa_d d_t$.

However, real interest rates will only partially offset the supply shock $\phi_s s_t$ if α is sufficiently large (i.e., the Fed places considerable weight on stabilizing the unemployment gap). In the limit, as $\alpha \to \infty$, the central bank does not react to the PC shock at all. This is

intuitive: the demand shock increases inflation and the output gap, creating no tradeoff. In contrast, the supply shock creates a tradeoff by increasing inflation while decreasing the output gap.

The equilibrium equation also provides insights into the magnitude of the monetary policy response. Suppose both demand and supply shocks are equally persistent ($\rho_d = \rho_s = \rho > 0$) and have an identical direct effect on inflation and on the output gap ($\kappa_d = \phi_s = 1$). Then, under standard parameter values, the central bank will respond more aggressively to demand shocks than to supply shocks—a result consistent with standard macroeconomic intuition. To see this, note that $\Upsilon_d > \Upsilon_s$ if, and only if:

$$\lambda^2 + \alpha(1 - \beta\rho) - \lambda(1 - \rho) > 0 \tag{41}$$

A sufficient condition for the inequality is $\lambda > 1 - \rho$. This holds with $\lambda = 0.3$ and $\rho > 0.7$, values consistent with our calibration, drawn from Clarida, Galí, and Gertler (2000) and Debortoli et al. (2018). Thus, under usual assumptions:

$$0 < \Upsilon_s < \Upsilon_d \tag{42}$$

In summary, holding all else constant, (i) the central bank's response to the demand factor is unambiguous, while its response to the supply factor involves a trade-off; and (ii) if supply and demand shocks are symmetric in persistence and have identical factor loadings, monetary policy will systematically show greater sensitivity toward demand shocks. After the central bank sets the optimal interest rates, the true realization of inflation and of the output gap will be given by the true realization of the innovations of d_t and s_t.

B.2 Asset Prices

We now augment the New Keynesian model with a risky asset and nominal risk-free bonds. We follow the asset-pricing approach in Cieslak and Pflueger (2023) and Pflueger and Rinaldi (2022). Aggregate consumption equals output, $C_t = Y_t$, and $c_t = \log C_t$. Recall that $y_t = x_t + z_t$, where x_t is the output gap and z_t denotes potential output.

Under the optimal discretionary policy characterized above, equilibrium inflation and the output gap depend only on the perceived supply factor s_t:

$$\pi_t = k_\pi^s s_t, \tag{43}$$

$$x_t = k_x^s s_t, \tag{44}$$

with

$$k_\pi^s = \frac{\alpha\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} > 0, \tag{45}$$

$$k_x^s = -\frac{\lambda\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} < 0. \tag{46}$$

Demand shocks are fully neutralized by optimal policy and do not enter equilibrium π_t or x_t.

As in CGG, we assume that a positive supply shock reduces potential output and output to the same extent, ceteris paribus. Specifically,

$$z_t = \bar{z} - z_s s_t, \qquad z_s > 0, \tag{47}$$

so that output will be

$$y_t = x_t + z_t = \bar{z} + y_s s_t, \qquad y_s := k_x^s - z_s < 0. \tag{48}$$

In equilibrium, a supply shock lowers potential output, and the output gap contracts accordingly. Because the gap narrows only if actual output falls with potential, output itself must decline.

Preferences are CRRA:

$$U_0 = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma}, \tag{49}$$

with risk aversion $\gamma > 0$ and discount factor $\beta \in (0,1)$. The log stochastic discount factor (SDF) is

$$m_{t+1} = \log\beta - \gamma(y_{t+1} - y_t) = \log\beta - \gamma y_s(s_{t+1} - s_t). \tag{50}$$

with the last equality holding in equilibrium. Conditional on information at time t, this can be written as

$$m_{t+1} = \tilde{m}_t - \gamma y_s s_{t+1}, \tag{51}$$

where $\tilde{m}_t$ is deterministic at time t. Since s_{t+1} is conditionally normal, $(m_{t+1}, \pi_{t+1}, \mathbb{E}_{t+1}\pi_{t+2})$ are jointly normal as in Cieslak and Pflueger (2023).

Nominal bonds. There are nominal risk-free bonds paying one dollar in one and two periods, with prices $\exp(-i_t)$ and $\exp(-2i_t^{(2)})$, respectively. There is also a real one-period risk-free bond with price $\exp(-r_t)$. To focus on inflation risk, we take the real rate r_t as approximately constant over the horizon of interest.

For the two-period nominal yield, we adopt the log-linear decomposition in Cieslak and Pflueger (2023). Under the assumption of an approximately constant real rate,

$$i_t^{(2)} = \frac{1}{2}\big(i_t + \mathbb{E}_t i_{t+1}\big) + \frac{1}{2}Cov_t\big(m_{t+1}, \pi_{t+1}^e\big) - \left[\frac{1}{2}Cov_t\big(\pi_{t+1}, \pi_{t+1}^e\big) + \frac{1}{4}Var_t\big(\pi_{t+1}^e\big)\right], \tag{52}$$

where

$$\pi_{t+1}^e := \mathbb{E}_{t+1}\pi_{t+2}.$$

The first term,

$$\frac{1}{2}\big(i_t + \mathbb{E}_t i_{t+1}\big),$$

is the expectations-hypothesis (EH) component: the average physical-measure expectation of the short rate over the bond's life. It is the yield implied by the EH in the absence of risk premia (equivalently, under risk neutrality), up to the usual Jensen/convexity adjustments. The second term,

$$\frac{1}{2}Cov_t\big(m_{t+1}, \pi_{t+1}^e\big),$$

is the inflation risk-premium. We focus on the covariance term because it provides a particularly transparent and intuitive decomposition of the bond risk premium into a priced inflation-risk channel. Following Cieslak and Pflueger (2023), we emphasize this channel as a tractable benchmark, while recognizing that empirical term premia may also reflect other sources of priced risk (e.g., real-rate risk) and convexity effects. The bracketed terms collect Jensen inequality corrections.

Using the equilibrium Taylor rule (Eq. (40)) and the AR(1) processes for (d_t, s_t), the risk-neutral component can be written as

$$\frac{1}{2}\big(i_t + \mathbb{E}_t i_{t+1}\big) = \theta_s s_t + \theta_d d_t, \tag{53}$$

with

$$\theta_s = \frac{1}{2}\left[\frac{\alpha\phi_s\rho_s + \frac{\lambda}{\psi}(1-\rho_s)\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)}\right](1+\rho_s), \tag{54}$$

$$\theta_d = \frac{1}{2}\frac{\kappa_d}{\psi}(1+\rho_d), \tag{55}$$

Under our calibration and in the standard New Keynesian model, we have $0 < \theta_s < \theta_d$, so demand shocks have a larger impact than supply shocks on the expected path of short rates.

To characterize the inflation risk premium, note that under the equilibrium dynamics,

$$\pi_{t+1} = k_\pi^s s_{t+1}, \tag{56}$$

$$\pi_{t+1}^e = \mathbb{E}_{t+1}\pi_{t+2} = k_\pi^s \rho_s s_{t+1}. \tag{57}$$

With $s_{t+1} = \rho_s s_t + \varepsilon_{s,t+1}$, we obtain

$$Cov_t(m_{t+1}, \pi_{t+1}^e) = Cov_t\left(\tilde{m}_t - \gamma y_s s_{t+1},\, k_\pi^s \rho_s s_{t+1}\right) = -\gamma y_s k_\pi^s \rho_s Var_t(\varepsilon_{s,t+1}). \tag{58}$$

Because $y_s < 0$, we have $-\gamma y_s k_\pi^s \rho_s > 0$, so the covariance between the SDF and expected future inflation is positive: higher expected inflation coincides with lower consumption and high marginal utility. Supply-driven inflation is therefore "bad" inflation, and nominal bonds are risky, or "stock-like" (Pflueger, 2025). In contrast, demand shocks affect only the risk-neutral path $\frac{1}{2}(i_t + \mathbb{E}_t i_{t+1})$ through d_t and do not enter π_t or y_t. In this baseline model, demand-driven inflation is completely stabilized and carries no risk premium.

Risky asset. We model the risky asset as in Pflueger and Rinaldi (2022). The risky asset is a claim to the future consumption stream, $C_{t+h}, h = 1, \ldots, \infty$, sold at price P_t^c. There is an aggregate levered firm that buys the consumption claim for P_t^c at t and sells equity worth θP_t^c, while $(1-\theta)P_t^c$ is financed by one-period risk-free debt. For simplicity, we focus on a one-period, zero-coupon consumption claim, with price:

$$P_{1,t}^c = \mathbb{E}_t[M_{t+1}C_{t+1}] \tag{59}$$

and its (log) return is approximately

$$r_{t+1}^c - r_t \approx (y_{t+1} - y_t) = y_s(s_{t+1} - s_t).$$

The standard log-linear approximation implies

$$\mathbb{E}_t[r^c_{t+1} - r_t] + \frac{1}{2} Var_t(r^c_{t+1}) = -Cov_t(m_{t+1}, r^c_{t+1}). \tag{60}$$

Using $m_{t+1} = \tilde{m}_t - \gamma y_s s_{t+1}$ and r_t constant, we obtain

$$-Cov_t(m_{t+1}, r^c_{t+1}) = \gamma y_s^2 Var_t(s_{t+1}) > 0, \tag{61}$$

so the consumption claim commands a positive risk premium. Absent supply shocks, consumption would be perfectly stabilized and the claim would carry no risk premium.

State-dependent volatility and time-varying premia. In the homoskedastic AR(1) benchmark, the perceived supply shock satisfies $s_{t+1} = \rho_s s_t + \varepsilon_{s,t+1}$ with $Var_t(\varepsilon_{s,t+1}) = \sigma_s^2$, so that $Var_t(s_{t+1}) = \sigma_s^2$ is constant. In this case the conditional variance of the pricing kernel and the relevant covariances with asset payoffs are constant, and bond and stock risk premia do not vary with the realization of s_t, as in standard New Keynesian models with homoskedastic shocks. To introduce state dependence in a parsimonious way, we assume that the conditional variance of perceived supply innovations follows

$$Var_t(\varepsilon_{s,t+h}) \;=\; \mu_s + \tau_s^h s_t^2, \qquad h = 1, 2, \ldots. \tag{62}$$

Under this reduced-form specification, large adverse supply shocks not only lower output and raise inflation, but also increase uncertainty about future inflation and activity. The resulting rise in $Var_t(m_{t+1})$ and in the covariances of payoffs with the pricing kernel makes both long nominal bonds and the consumption claim especially risky in those states, so risk premia increase in s_t^2.

This assumption is made for tractability. One can also obtain endogenously time-varying risk premia by embedding habit formation into the New Keynesian environment, which makes the SDF's exposure to homoskedastic shocks state dependent (Campbell, Pflueger, and Viceira, 2020; Pflueger and Rinaldi, 2022; Cieslak and Pflueger, 2023; Pflueger, 2025). Such extensions introduce additional state variables and parameters that are orthogonal to the identification questions emphasized here. The reduced-form specification in (62) therefore provides a simple way to capture the same economic intuition while preserving analytical transparency. The reduced-form specification captures the widely documented fact that macroeconomic and inflation uncertainty rises in adverse states, such as recessions or large negative supply shocks (Fernández-Villaverde et al., 2011; Basu and Bundick, 2017).

We calibrate μ_s and τ_s using the moment restrictions implied by Equation (62). First, the unconditional variance of $\varepsilon_{s,t}$ matches the empirical second moment of the estimated supply factor. Second, we assume that in a "two-sigma" supply state $s_t = 2\sqrt{\mathbb{E}[s_t^2]}$, the conditional variance $\text{Var}_t(\varepsilon_{s,t+1})$ is twice its unconditional level. These conditions imply the linear system

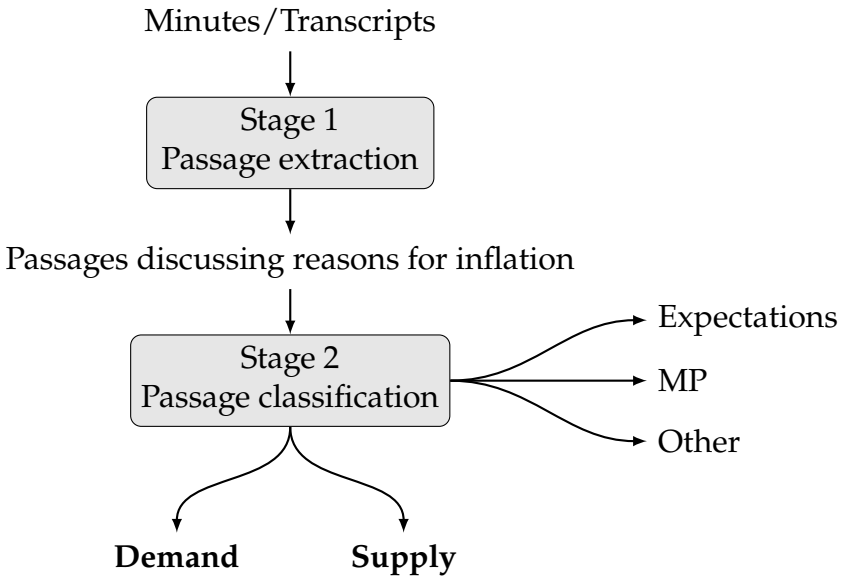
$$\mathbb{E}[Var_t(\varepsilon_{s,t+1})] = \mu_s + \tau_s \mathbb{E}[s_t^2] \tag{63}$$

$$2\mathbb{E}[Var_t(\varepsilon_{s,t+1})] = \mu_s + 4\tau_s \mathbb{E}[s_t^2] \tag{64}$$

which yields $\tau_s = \frac{\mathbb{E}[Var_t(\varepsilon_{s,t+1})]}{3\mathbb{E}[s_t^2]}$ and $\mu_s = \frac{2}{3}\mathbb{E}[Var_t(\varepsilon_{s,t+1})]$. We implement this calibration using the empirical analogs of these moments, as in stochastic-volatility calibrations in macro-finance, which pin down volatility parameters by matching unconditional and state-contingent variances of macro variables or asset returns (Fernández-Villaverde et al., 2011).

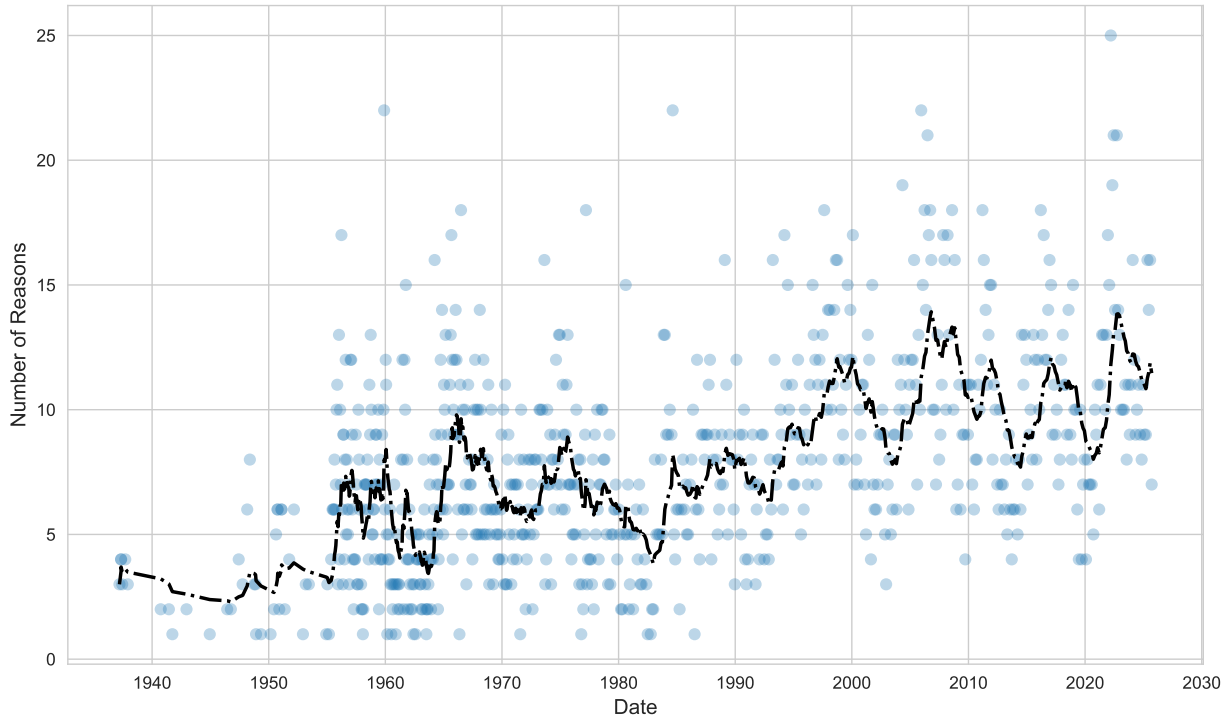
C Supplementary Figures

Figure C.1: Diagram for two-stage reason classification procedure



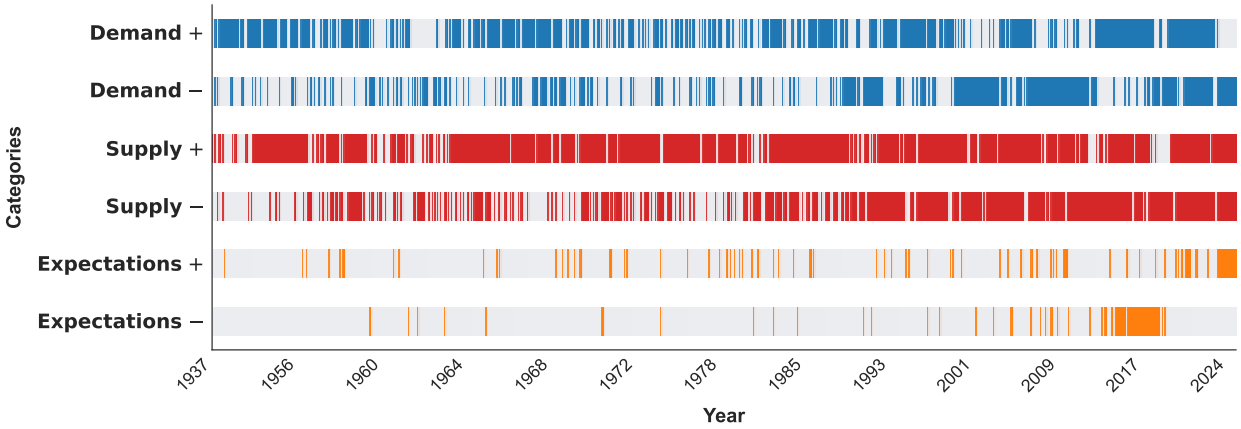
Notes: The diagram summarizes our two-stage procedure for extracting inflation reasons from FOMC textual records. Stage 1 identifies passages in the minutes where participants cite structural drivers of inflation. Stage 2 classifies reasons as demand, supply, expectations, monetary policy, or other. See the main text for more details.

Figure C.2: Total Reason Counts Over Time



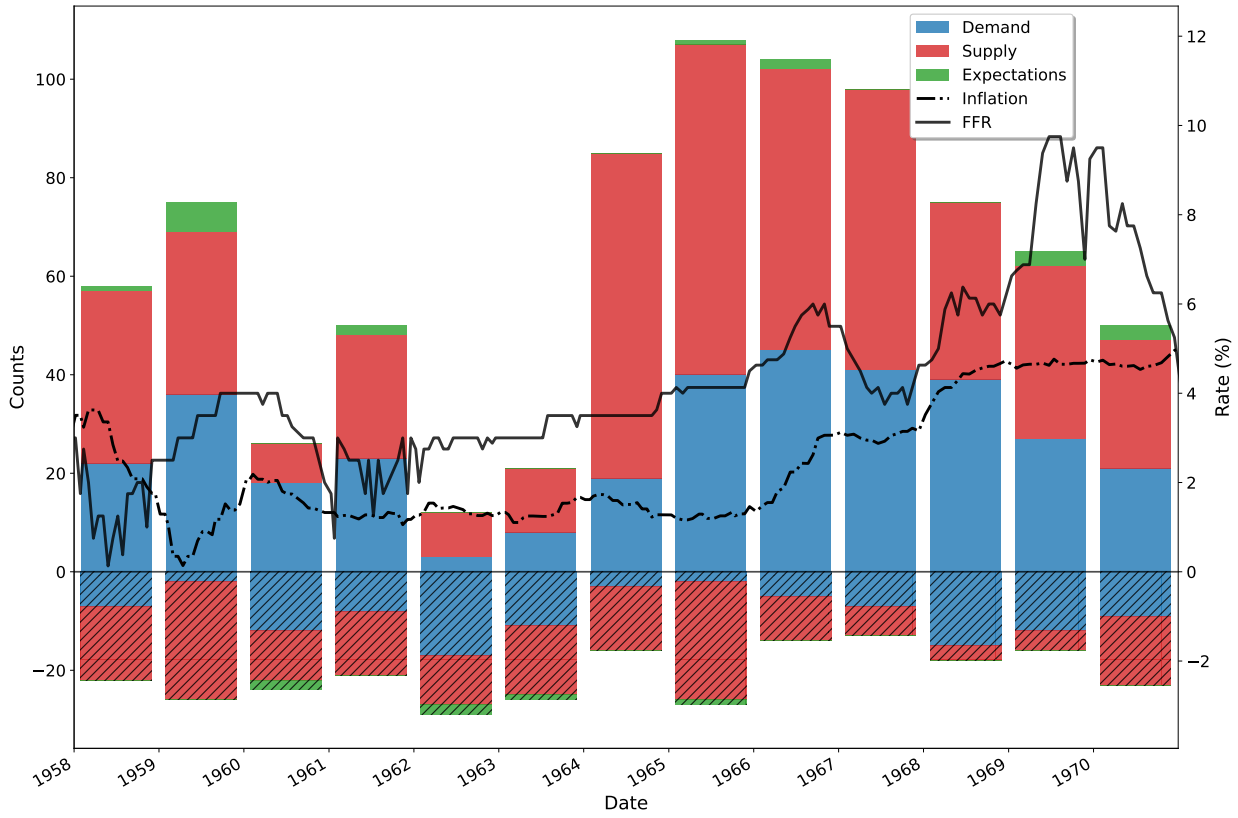
Note: Each dot gives the meeting-level count of reasons—demand, supply, or expectations, inflationary or disinflationary. The black curve is an exponential moving average with a half-life of eight meetings. See the main text for details on the construction of these variables.

Figure C.3: Presence of demand, supply, and expectations reasons over Time



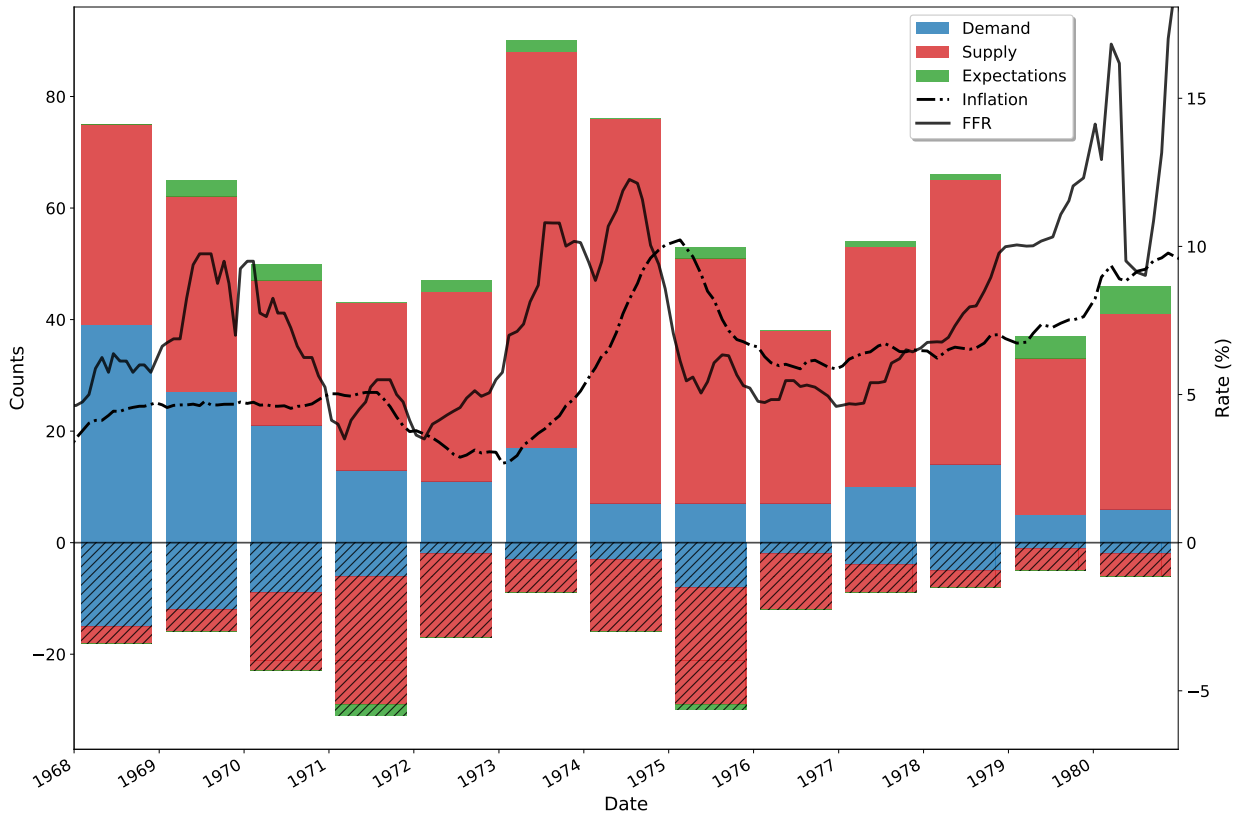
Note: A filled cell marks any FOMC meeting which mentions at least one reason of the indicated category. Inflationary (disinflationary) reasons are accompanied by a "+" ("−"). See the main text for details on the construction of the variables.

Figure C.4: Aggregated Reason Counts Over Time (1960s)



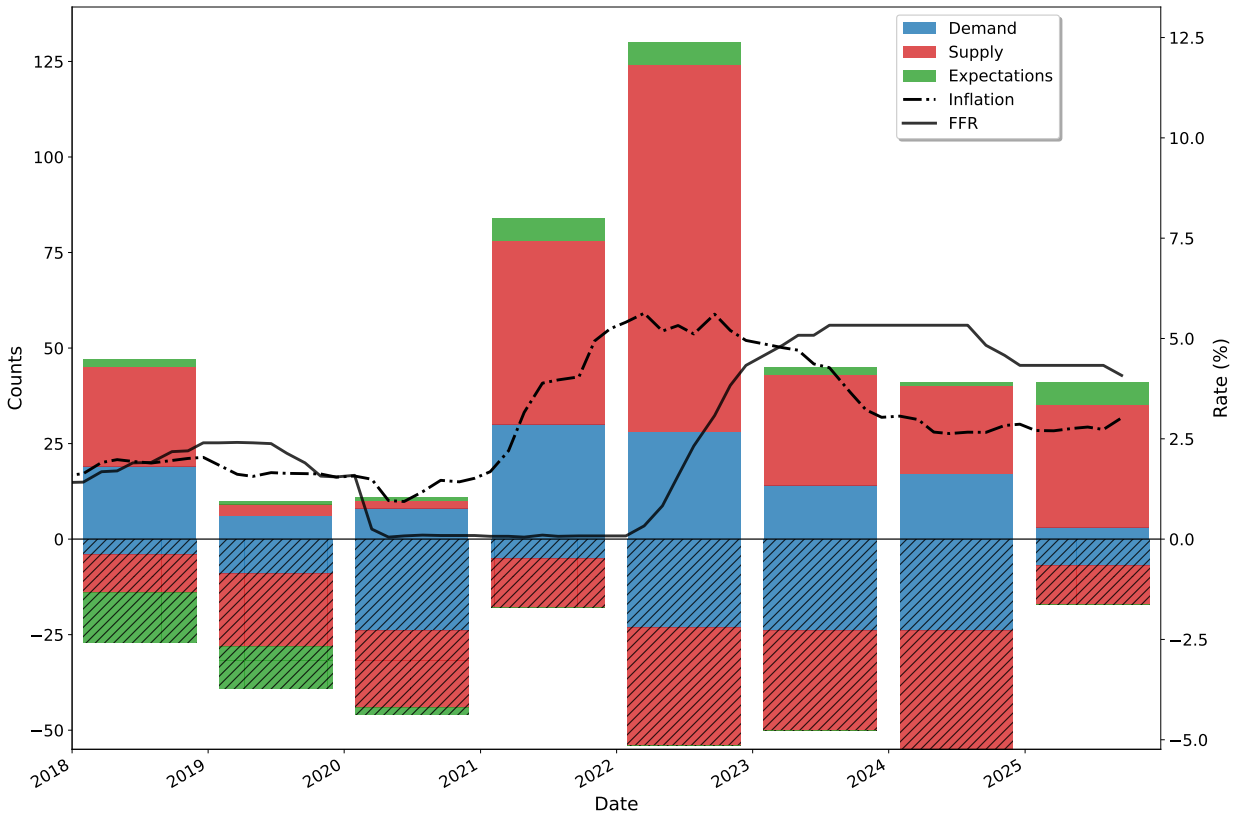
Notes: This figure aggregates annual counts of demand, supply, and expectation reasons, distinguishing inflationary (bars going up, solid) and disinflationary reasons (bars going down, hatched) on the left axis. The snapshot covers the 1960s. The Federal Funds Rate and inflation are shown on the right axis. We use CPI inflation before 1960 and core inflation afterwards. See the main text for details on variable construction.

Figure C.5: Aggregated Reason Counts Over Time (1970s)



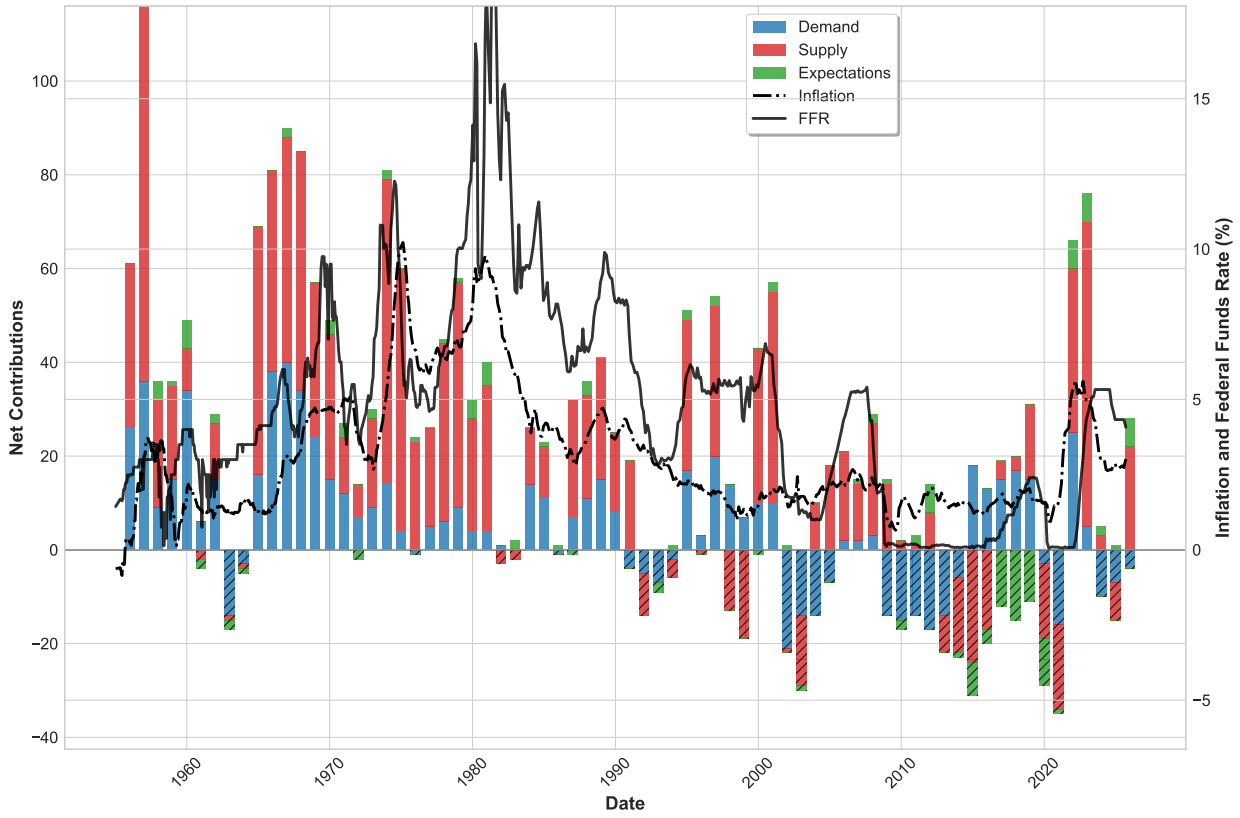
Notes: This figure aggregates annual counts of demand, supply, and expectation reasons, distinguishing inflationary (solid) and disinflationary (hatched) bars on the left axis. The snapshot covers the 1970s. The Federal Funds Rate and core inflation are shown on the right axis. See the main text for details on variable construction.

Figure C.6: Aggregated Reason Counts Over Time (Covid)



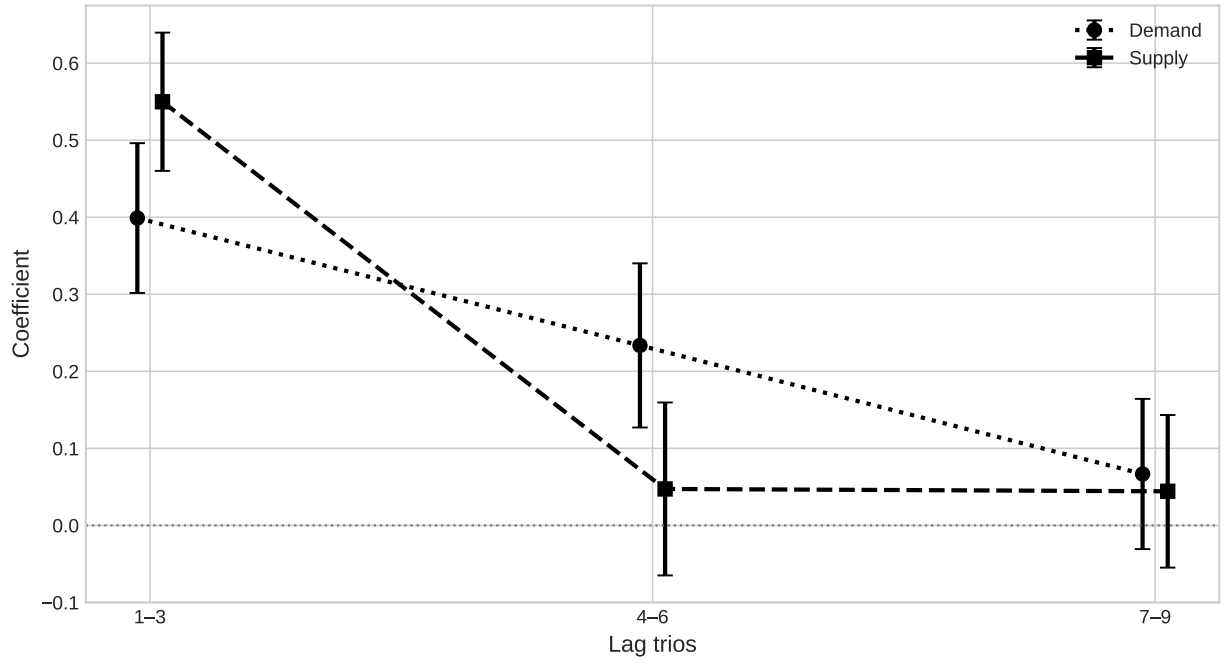
Notes: This figure aggregates annual counts of demand, supply, and expectation reasons, separating inflationary (solid) and disinflationary (hatched) bars on the left axis. The snapshot covers the Covid period. The Federal Funds Rate and core inflation are plotted on the right axis. See the main text for details on variable construction.

Figure C.7: Net Aggregated Reason Counts Over Time



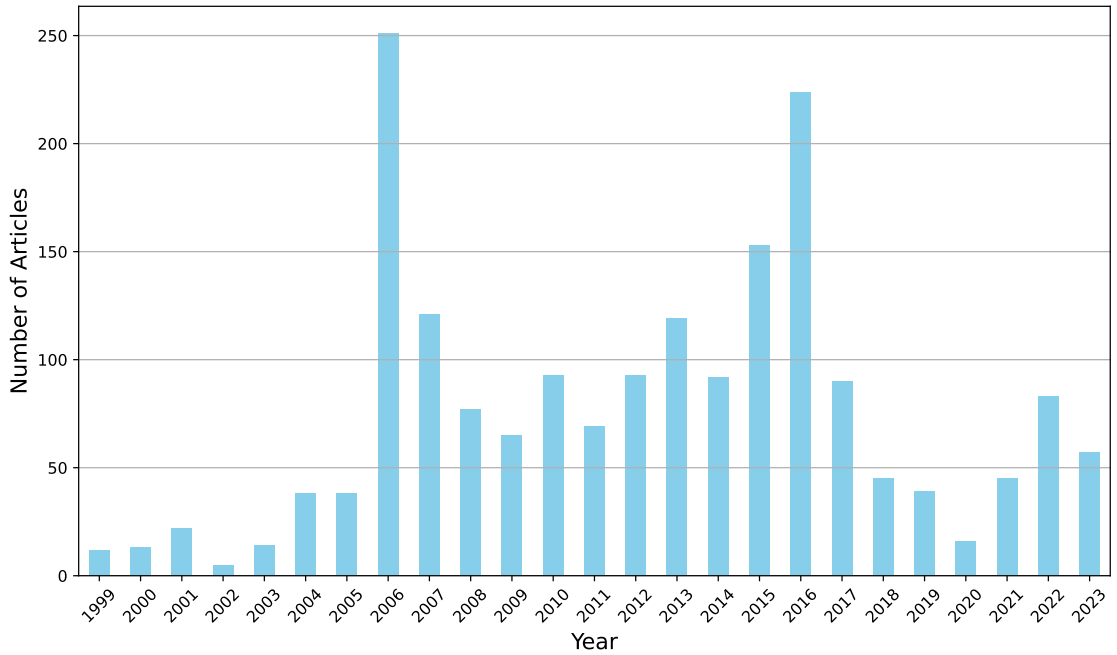
Notes: The figure plots annual net reason counts for demand, supply, and expectations on the left axis. Net imbalance equals the number of reasons indicating higher inflation minus those indicating lower inflation in a given year and category. The right axis shows the Federal Funds Rate and inflation. We use CPI inflation before 1960 and core inflation thereafter. See the main text for variable construction details.

Figure C.8: Persistence of imbalances



Notes: This figure plots autoregressive coefficients with the associated 95% confidence intervals. Each imbalance is related to the average of adjacent lag trios, estimated separately for demand and supply. Specifically, we regress contemporaneous imbalance on theree subsequent averages: averages of lags one to three, averages of lags four to six, and averages of lags seven to nine. We use Newey–West standard errors with 8 lags. See the main text for details of the variable construction.

Figure C.9: Number of newspaper articles per year



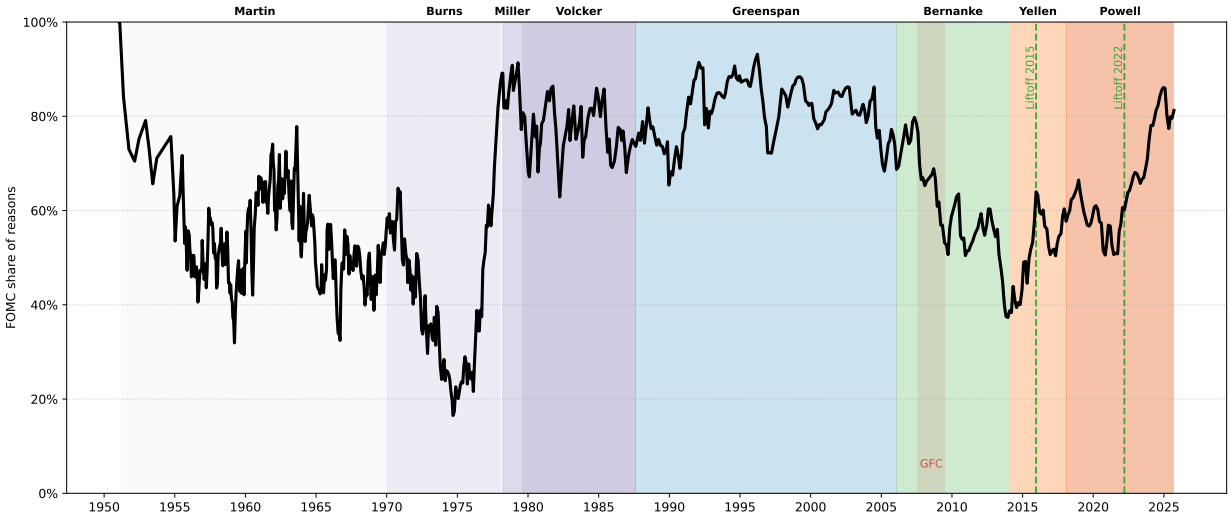
Notes: This plot presents the number of newspaper articles considered in our sample per year.

Figure C.10: Wordcloud of the body of the newspaper articles considered



Notes: The figure presents a wordcloud of from the initial filtered set of articles we consider. We lemmatize and filter tokens, removing stopwords, punctuation, and numbers. The stopword list combines the standard lists with additional domain-specific terms—both generic media boilerplate and FOMC-specific terms. Only nouns, proper nouns, and adjectives are retained.

Figure C.11: Share of reasons discussed by staff and FOMC members over time



Notes: This figure plots an exponential moving average of the share of reasons presented by FOMC members in each meeting over the sample period, using a 4-meeting half-life. Shaded regions identify the tenure of each FOMC Chair.

D Supplementary Tables

Table D.1: Sample Periods and Number of Observations by Series

Series	Start	End	Obs.
FOMC Records			
Full corpus	Jan 1937	Jun 2025	789
Largest estimation sample	Jan 1960	Jun 2025	636
Statement imbalances	Aug 1994	Dec 2023	202
News articles imbalances	May 1999	Jun 2023	161
Macroeconomic series			
PCE Core Inflation	Jan 1960	Jun 2025	636
PCE Headline Inflation	Jan 1960	Jun 2025	636
Federal Funds Effective Rate	Sep 1954	Jun 2025	723
Shadow Rate	Feb 1990	Jan 2022	248
Monetary policy shock variables			
Macroeconomic News	Feb 1988	Dec 2023	361
Yield series			
Risk-neutral yield	June 1961	Jun 2025	616
Term premium	June 1961	Jun 2025	616

Notes: This table reports the start date, end date, and total number of observations for each series used in the empirical analysis. See Section 1 for more details about the data sources.

Table D.2: Top reasons for inflation by category

Demand		**Supply**		**Expectations**	
Reason	Count	Reason	Count	Reason	Count
resource slack	17	energy prices	66	lower longer-term inflation expectations	12
high resource utilization	16	dollar depreciation	40	rising inflation expectations	5
slack in labor and product markets	14	productivity gains	26	low inflation expectations	3
slack in resource utilization	14	wage increases	21	declining inflation expectations	3
economic slack	9	oil prices	19	energy prices raising inflation expectations	2
tight labor markets	8	food prices	19	declines in market-based inflation compensation	2
labor market strengthening	7	unit labor costs	18	low inflation compensation	2
underutilized resources (slack)	7	commodity prices	16	oil prices	2
tightening resource utilization	7	steel price increases	13	upward drift in inflation expectations	2
tight labor market	6	higher import prices	13	worsening inflation expectations	2
Total	**105**	**Total**	**251**	**Total**	**35**

Notes: The table reports the ten most cited explanations for inflation in the minutes, sorted by category. "Demand'', "Supply'', and "Expectations" reflect the model's tags. "Count'' is the number of occurrences across the full sample. See the main text for details on the construction of these variables.

Table D.3: Yields and Term Premia on imbalances – Post GFC

	1y	2y	5y	10y	1y	2y	5y	10y
I_s	0.042	0.035	0.033	0.032	0.110***	0.156***	0.216***	0.289***
	(0.040)	(0.044)	(0.044)	(0.037)	(0.024)	(0.033)	(0.055)	(0.071)
I_d	0.070	0.115	0.132*	0.109**	-0.059**	-0.053	-0.067	-0.149
	(0.067)	(0.070)	(0.068)	(0.055)	(0.028)	(0.039)	(0.066)	(0.092)
π_t	0.193***	0.187***	0.157***	0.120***	-0.016	-0.035	-0.075***	-0.100***
	(0.064)	(0.049)	(0.036)	(0.027)	(0.017)	(0.023)	(0.027)	(0.037)
x_t	0.028	0.029	0.030	0.025	-0.007	-0.007	-0.004	0.002
	(0.027)	(0.025)	(0.022)	(0.017)	(0.011)	(0.014)	(0.020)	(0.024)
$rn_{t-2}^{(n)}$	0.901***	0.887***	0.868***	0.861***				
	(0.045)	(0.043)	(0.043)	(0.045)				
$tp_{t-2}^{(n)}$					0.722***	0.784***	0.831***	0.844***
					(0.099)	(0.073)	(0.054)	(0.052)
Obs.	137	137	137	137	137	137	137	137
R^2	0.969	0.964	0.954	0.949	0.755	0.807	0.843	0.855

Notes: Each column reports OLS estimates of risk-neutral yields $rn_t^{(n)}$ or term premia $tp_t^{(n)}$ on demand and supply imbalances (I_d, I_s), core inflation (π_t), the log output gap (x_t), and the dependent variable lagged by two meetings, analogous to the auxiliary regression in Equation (18). The sample is restricted to 2008 - 2025. Maturities are one, two, five, and ten years. Risk-neutral yields and term premia come from Adrian, Crump, and Moench (2013); yields and premia are measured two business days after FOMC meeting t. The sample spans August 1961–June 2025 at the FOMC-meeting frequency. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table D.4: Augmented Generalized Taylor Rule Estimates by Era – FOMC and Staff

	FOMC				*Staff*			
	Full	*Pre-Volcker*	*Volcker–GFC*	*Post-GFC*	*Full*	*Pre-Volcker*	*Volcker–GFC*	*Post-GFC*
ϕ_π	1.38***	0.77***	1.81***	3.06*	1.43***	0.7**	1.83***	2.98*
	(0.34)	(0.27)	(0.19)	(1.83)	(0.34)	(0.27)	(0.2)	(1.69)
ϕ_x	1.06***	1.16***	0.66***	0.37	1.17***	1.04***	0.72***	0.47
	(0.26)	(0.37)	(0.17)	(0.37)	(0.26)	(0.34)	(0.2)	(0.32)
ϕ_s	0.18	-0.02	-0.03	0.17	-0.32	0.89	-0.31	-1.45
	(0.44)	(0.66)	(0.3)	(0.98)	(0.49)	(0.65)	(0.47)	(1.43)
ϕ_d	1.51**	-1.03	1.75***	1.46	0.28	-0.28	1.11***	0.58
	(0.74)	(0.85)	(0.4)	(1.24)	(0.57)	(0.86)	(0.28)	(0.99)
ρ	0.87***	0.88***	0.63***	0.92***	0.87***	0.88***	0.62***	0.92***
	(0.03)	(0.03)	(0.07)	(0.04)	(0.03)	(0.04)	(0.08)	(0.04)
N	638	270	229	139	638	270	229	139
R^2	0.931	0.919	0.901	0.959	0.931	0.919	0.898	0.960

Notes: Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (17), for different samples: *Full sample* = 1960Q1–2025Q2; *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. We estimate the structural slopes using imbalances constructed with reasons expressed by FOMC participants and by staff separately. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. φ_π is the response to π_t (core inflation), φ_x to x_t (log output gap); φ_d and φ_s to the demand and supply textual imbalances. ρ is the interest-rate smoothing parameter. Eq. (18) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table D.5: Direct estimates of policy rate and imbalances

	Full sample	*Pre-Volcker*	*Volcker–GFC*	*Post-GFC*
$I_{d,t}$	0.298***	0.130	0.677***	0.209**
	(0.093)	(0.087)	(0.238)	(0.106)
$I_{s,t}$	0.174***	0.173***	0.146	0.107
	(0.054)	(0.059)	(0.095)	(0.139)
ρ	0.951***	0.962***	0.904***	0.978***
	(0.022)	(0.032)	(0.036)	(0.031)
Obs.	723	353	229	141
R^2	0.925	0.922	0.863	0.943

Notes: Each row reports reduced-form slope coefficients from the estimate of the policy rate on past quarter policy rate and demand and supply imbbalances. We estimate the regression separately for different samples: *Full sample* = 1960Q1–2025Q2; *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. ρ is the interest-rate smoothing parameter. We construct Newey-West standard errors (with 8 lags). *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table D.6: Regression on imbalances from statements on imbalances from minutes and news articles

	$I^{stat}_{d,t}$		$I^{stat}_{s,t}$	
$I_{s,t}$	0.101		0.257***	
	(0.118)		(0.072)	
$I_{d,t}$	0.597***		-0.140	
	(0.138)		(0.088)	
$I^{news}_{s,t}$		-0.041		0.481***
		(0.208)		(0.150)
$I^{news}_{d,t}$		0.352***		0.045
		(0.070)		(0.043)
Obs.	191	151	191	151
R^2	0.197	0.188	0.114	0.127

Notes: The table presents regression results of demand and supply imbalances constructed from the statements ($I^{stat}_{d,t}$ and $I^{stat}_{s,t}$) on the corresponding imbalances constructed previously using minutes and transcripts ($I_{d,t}$ and $I_{s,t}$) and news articles ($I^{news}_{d,t}$ and $I^{news}_{s,t}$). See the main text for more details of the construction of these variables. We construct Newey-West standard errors (with 8 lags). *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table D.7: Regressions of high-frequency asset prices changes on the difference between statament and news imbalances

	MPS	*S&P 500*	*ED1*	*ED4*	*VIXY*
$I^{stat}_{d,t} - I^{news}_{d,t}$	-0.001	-0.034	0.003	-0.004	-0.002
	(0.004)	(0.044)	(0.004)	(0.004)	(0.002)
$I^{stat}_{s,t} - I^{news}_{s,t}$	0.013	-0.198***	0.015*	0.008	0.006***
	(0.009)	(0.073)	(0.009)	(0.009)	(0.002)
Obs.	151	151	151	151	87
R^2	0.210	0.057	0.142	0.211	0.120
Controls	✓	✓	✓	✓	✓

Notes: The table reports regressions of high-frequency changes in asset prices and the monetary policy shock of Bauer and Swanson (2023) on the difference between demand and supply imbalances extracted from FOMC statement and the corresponding imbalances extracted from news articles. Asset price changes are from ten minutes before to twenty minutes after the announcement. The dependent variables include the change in the S&P 500 e-mini futures, one- and two-quarter ahead Eurodollar futures (ED1 and ED4), VIXY, and the monetary policy shock from Bauer and Swanson (2023). VIXY is an exchange-traded fund tracking short-term VIX futures; changes in VIXY are computed from TAQ data. The control set corresponds to the macro news predictors in Bauer and Swanson (2023), described in Section 4.6. The details about the news articles are given in Appendix A.5. Sample runs from May, 1999, to June, 2023. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table D.8: Regressions of high-frequency asset prices changes on statament and news imbalances

	MPS	*S&P 500*	*ED1*	*ED4*	*VIXY*
$I^{stat}_{d,t}$	0.001	-0.052	0.001	-0.001	-0.000
	(0.006)	(0.068)	(0.004)	(0.007)	(0.003)
$I^{stat}_{s,t}$	0.016*	-0.225***	0.017*	0.011	0.009***
	(0.009)	(0.074)	(0.009)	(0.010)	(0.002)
$I^{news}_{d,t}$	0.001	0.032	-0.004	0.005	0.002
	(0.004)	(0.049)	(0.004)	(0.005)	(0.002)
$I^{news}_{s,t}$	-0.005	0.105	-0.008	0.001	0.002
	(0.011)	(0.114)	(0.010)	(0.013)	(0.005)
Obs.	151	151	151	151	87
R^2	0.214	0.062	0.147	0.216	0.156

Notes: The table reports regressions of high-frequency changes in asset prices and the monetary policy shock of Bauer and Swanson (2023) on the demand and supply imbalances extracted from FOMC statement, controlling for the corresponding imbalances extracted from news articles. Asset price changes are from ten minutes before to twenty minutes after the announcement. The dependent variables include the change in the S&P 500 e-mini futures, one- and two-quarter ahead Eurodollar futures (ED1 and ED4), VIXY, and the monetary policy shock from Bauer and Swanson (2023). VIXY is an exchange-traded fund tracking short-term VIX futures; changes in VIXY are computed from TAQ data. The control set corresponds to the macro news predictors in Bauer and Swanson (2023), described in Section 4.6. The details about the news articles are given in Appendix A.5. Sample runs from May, 1999, to June, 2023. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table D.9: Regressions of inflation on imbalances (moving average)

	All		**FOMC**		**Staff**	
	Core	*Headline*	*Core*	*Headline*	*Core*	*Headline*
I_d^{ma}	0.052	0.077	0.073**	0.093*	0.040**	0.061*
	(0.038)	(0.062)	(0.029)	(0.048)	(0.020)	(0.033)
I_s^{ma}	0.180***	0.345***	0.095***	0.178***	0.078***	0.143***
	(0.040)	(0.067)	(0.030)	(0.052)	(0.020)	(0.033)
π_{t-2}	0.979***		0.981***		0.984***	
	(0.017)		(0.018)		(0.018)	
π_{t-2}		0.958***		0.964***		0.968***
		(0.021)		(0.023)		(0.021)
Obs.	634	634	634	634	634	634
R^2	0.979	0.957	0.978	0.955	0.978	0.955

Notes: This table replicates Table 2 using an alternative imbalance specification that scales net counts by a moving average of total counts over the last eight meetings. The dependent variables are monthly headline and core inflation from FRED, expressed as year-over-year percentage changes. Demand and supply imbalances are defined as in Appendix F.1. Regressions are estimated by OLS with Newey–West standard errors (8 lags, in parentheses). *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table D.10: Augmented Taylor Rule Estimates (moving average)

	All			
	Full sample	Pre-Volcker	Volcker–GFC	Post-GFC
ϕ_π	1.42***	0.69**	1.86***	3.32
	(0.36)	(0.27)	(0.17)	(2.04)
ϕ_x	1.14***	1.08***	0.76***	0.34
	(0.25)	(0.33)	(0.16)	(0.4)
ϕ_s	-0.26	1.33	-1.33**	-0.79
	(0.86)	(0.83)	(0.63)	(2.2)
ϕ_d	0.91	-1.16	3.18***	2.78
	(1.1)	(0.85)	(1.06)	(2.89)
ρ	0.87***	0.88***	0.6***	0.92***
	(0.03)	(0.03)	(0.08)	(0.04)
N	634	266	229	139
R^2	0.931	0.919	0.903	0.959

Notes: This table replicates Table 3 using an alternative imbalance specification that scales net counts by a moving average of total counts over the last eight meetings. Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (17), for different samples: *Full sample* = 1960Q1–2025Q2; *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. φ_π is the response to π_t (core inflation), φ_x to x_t (log output gap); φ_d and φ_s to the demand and supply textual imbalances. ρ is the interest-rate smoothing parameter. Demand and supply imbalances are defined as in Appendix F.1. Eq. (18) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table D.11: Augmented Generalized Taylor Rule Estimates by Era (All Sources)

	All			
	Full sample	Pre-Volcker	Volcker–GFC	Post-GFC
ϕ_π	1.4***	0.69**	1.85***	3.37**
	(0.36)	(0.27)	(0.18)	(1.67)
ϕ_x	1.1***	1.05***	0.74***	0.35
	(0.25)	(0.32)	(0.15)	(0.32)
ϕ_s	0.04	1.31	-0.99	-1.7
	(0.8)	(0.91)	(0.6)	(2.21)
ϕ_d	1.54	-0.65	3.41***	2.4
	(1.08)	(0.89)	(0.71)	(2.69)
ρ	0.87***	0.88***	0.61***	0.92***
	(0.03)	(0.03)	(0.08)	(0.04)
N	634	266	229	139
R^2	0.931	0.919	0.903	0.959

Table D.12: Interest rates, yields, and term premia's response controlling for macro news (moving average)

	FFR/ Shadow Rate	*Risk Neutral Yields*		*Term Premia*	
	i_t^s	1y	5y	1y	5y
I_s^{ma}	0.137**	0.095**	0.068	0.058***	0.072***
	(0.069)	(0.046)	(0.050)	(0.022)	(0.027)
I_d^{ma}	0.244	0.352***	0.354***	-0.007	0.046
	(0.161)	(0.115)	(0.117)	(0.034)	(0.041)
Obs.	284	284	284	284	284
R^2	0.982	0.978	0.964	0.884	0.971
Controls	✓	✓	✓	✓	✓

Notes: This table replicates Table 7 using an alternative imbalance specification that scales net counts by a moving average of total counts over the last eight meetings. The table re-estimates regressions of the policy rate, risk-neutral yields, and term premia on demand and supply imbalances, with and without the set of macro-news predictors of Bauer and Swanson (2023). Demand and supply imbalances are defined as in Appendix F.1. Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and the level of the dependent variable at $t-2$. See the main text for details on the variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table D.13: Perceptions of the Fed, Policy Rates, and the Yield Curve

	FFR/ Shadow Rate	*Risk Neutral Yields*		*Term Premia*	
	i_t^s	1y	5y	1y	5y
I_s	0.131*	0.099*	0.074	0.053***	0.068**
	(0.072)	(0.051)	(0.053)	(0.020)	(0.026)
I_d	0.270	0.310**	0.316**	0.000	0.039
	(0.166)	(0.139)	(0.133)	(0.036)	(0.046)
Obs.	284	284	284	284	284
R^2	0.983	0.978	0.964	0.884	0.970

Table D.14: Regressions of inflation on imbalances (inverse hyperbolic sine)

	All		**FOMC**		**Staff**	
	Core	*Headline*	*Core*	*Headline*	*Core*	*Headline*
I_d	0.022*	0.040**	0.035**	0.058**	0.019	0.048*
	(0.012)	(0.019)	(0.014)	(0.024)	(0.018)	(0.028)
I_s	0.056***	0.118***	0.043***	0.098***	0.079***	0.149***
	(0.013)	(0.020)	(0.014)	(0.027)	(0.021)	(0.028)
π_{t-2}	0.981***		0.983***		0.984***	
	(0.017)		(0.019)		(0.017)	
π_{t-2}		0.960***		0.966***		0.966***
		(0.020)		(0.023)		(0.020)
Obs.	634	634	634	634	634	634
R^2	0.979	0.959	0.978	0.956	0.979	0.957

Notes: This table replicates Table 2 using an alternative imbalance specification applies the inverse hyperbolic sine to the net counts. The dependent variables are monthly headline and core inflation from FRED, expressed as year-over-year percentage changes. Demand and supply imbalances are defined as in Appendix F.1. Regressions are estimated by OLS with Newey–West standard errors (8 lags, in parentheses). *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table D.15: Augmented Taylor Rule Estimates (inverse hyperbolic sine)

	All			
	Full sample	Pre-Volcker	Volcker–GFC	Post-GFC
ϕ_π	1.4***	0.67**	1.81***	3.69*
	(0.35)	(0.27)	(0.18)	(2.07)
ϕ_x	1.05***	1.02***	0.64***	0.29
	(0.25)	(0.33)	(0.15)	(0.38)
ϕ_s	-0.02	0.46	-0.18	-0.32
	(0.24)	(0.33)	(0.13)	(0.56)
ϕ_d	0.64**	-0.17	0.87***	0.81
	(0.32)	(0.35)	(0.14)	(0.66)
ρ	0.87***	0.88***	0.6***	0.93***
	(0.03)	(0.04)	(0.07)	(0.04)
N	634	266	229	139
R^2	0.931	0.919	0.903	0.959

Notes: This table replicates Table 3 using an alternative imbalance specification that applies the inverse hyperbolic sine on net counts. Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (17), for different samples: *Full sample* = 1960Q1–2025Q2; *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. φ_π is the response to π_t (core inflation), φ_x to x_t (log output gap); φ_d and φ_s to the demand and supply textual imbalances. ρ is the interest-rate smoothing parameter. Demand and supply imbalances are defined as in Appendix F.1. Eq. (18) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table D.16: Interest rates, yields, and term premia's response controlling for macro news (inverse hyperbolic sine)

	FFR/ Shadow Rate	*Risk Neutral Yields*		*Term Premia*	
	i_t^s	1y	5y	1y	5y
I_s^{asine}	0.037**	0.022	0.011	0.013**	0.021***
	(0.018)	(0.014)	(0.013)	(0.005)	(0.007)
I_d^{asine}	0.066**	0.077***	0.080***	0.001	0.011
	(0.033)	(0.030)	(0.028)	(0.007)	(0.008)
Obs.	284	284	284	284	284
R^2	0.983	0.978	0.965	0.884	0.971
Controls	✓	✓	✓	✓	✓

Notes: This table replicates Table 7 using an alternative imbalance specification that applies the inverse hyperbolic sine to net counts. The table re-estimates regressions of the policy rate, risk-neutral yields, and term premia on demand and supply imbalances, with and without the set of macro-news predictors of Bauer and Swanson (2023). Demand and supply imbalances are defined as in Appendix F.1. Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and the level of the dependent variable at $t-2$. See the main text for details on the variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table D.17: Regressions of inflation on imbalances

	All		**FOMC**		**Staff**	
	Core	*Headline*	*Core*	*Headline*	*Core*	*Headline*
I_d	0.064**	0.123***	0.079***	0.137***	0.047*	0.075*
	(0.026)	(0.045)	(0.026)	(0.046)	(0.026)	(0.044)
I_s	0.131***	0.261***	0.097***	0.199***	0.111***	0.233***
	(0.031)	(0.050)	(0.027)	(0.056)	(0.030)	(0.045)
π_{t-2}	0.981***	0.961***	0.982***	0.966***	0.985***	0.968***
	(0.017)	(0.020)	(0.019)	(0.023)	(0.018)	(0.020)
Obs.	630	630	630	630	630	630
R^2	0.979	0.957	0.978	0.955	0.978	0.956

Notes: The dependent variables are monthly levels of headline and core inflation, expressed as year-over-year percentage changes, from the FRED. Demand and supply imbalances are defined as in Equation (14). We compute imbalances separately for FOMC members, staff, and the pooled set of reasons. Regressions are estimated by OLS with Newey-West standard errors with 8 lags (in parentheses). *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table D.18: Interest rates, yields, and term premia's response controlling for macro news – FOMC members and staff

	FFR/ Shadow Rate	*Risk Neutral Yields*		*Term Premia*	
	i_t^s	1y	5y	1y	5y
Panel A: FOMC Reasons					
I_s	0.081**	0.042	0.016	0.022*	0.042**
	(0.034)	(0.031)	(0.032)	(0.013)	(0.017)
I_d	0.179**	0.206***	0.182***	0.009	0.019
	(0.075)	(0.053)	(0.051)	(0.015)	(0.021)
Obs.	284	284	284	284	284
R^2	0.983	0.978	0.964	0.882	0.970
Controls	✓	✓	✓	✓	✓
Panel B: Staff Reasons					
I_s	0.014	0.039	0.019	0.006	0.009
	(0.049)	(0.037)	(0.034)	(0.014)	(0.018)
I_d	0.064	0.103**	0.108**	−0.003	0.019
	(0.054)	(0.051)	(0.055)	(0.017)	(0.020)
Obs.	284	284	284	284	284
R^2	0.982	0.977	0.963	0.881	0.970
Controls	✓	✓	✓	✓	✓

Notes: The table re-estimates regressions of the policy rate, risk-neutral yields, and term premia on demand and supply imbalances, with and without the set of macro-news predictors of Bauer and Swanson (2023). We compute the regressions separately for imbalances constructed with FOMC and staff reasons. Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and the level of the dependent variable at $t-2$. See Appendix F.2 for details on the variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table D.19: Greenbook Forecasts and FOMC/Staff Imbalances

Panel A: Inflation Greenbook Forecasts					
	$h=0$	$h=1$	$h=2$	$h=3$	$h=4$
I_d^{FOMC}	0.146	0.232*	0.139**	0.076	0.053
	(0.221)	(0.139)	(0.066)	(0.049)	(0.035)
I_s^{FOMC}	0.703***	0.158*	0.133	0.102**	0.079**
	(0.265)	(0.094)	(0.089)	(0.045)	(0.034)
I_d^{Staff}	−0.035	−0.048	−0.035	0.023	0.035
	(0.218)	(0.107)	(0.063)	(0.038)	(0.034)
I_s^{Staff}	0.291*	0.170	0.079	0.091*	0.024
	(0.158)	(0.119)	(0.054)	(0.055)	(0.029)
Obs.	279	279	279	279	279
R^2	0.522	0.770	0.915	0.952	0.967
Panel B: Unemployment Greenbook Forecasts					
	$h=0$	$h=1$	$h=2$	$h=3$	$h=4$
I_d^{FOMC}	−0.188***	−0.160**	−0.139**	−0.139**	−0.134**
	(0.072)	(0.067)	(0.066)	(0.064)	(0.066)
I_s^{FOMC}	−0.001	0.004	0.023	0.020	0.033
	(0.033)	(0.036)	(0.037)	(0.038)	(0.040)
I_d^{Staff}	−0.109***	−0.128***	−0.140***	−0.126**	−0.122**
	(0.041)	(0.048)	(0.054)	(0.053)	(0.055)
I_s^{Staff}	0.038	0.037	0.034	0.033	0.006
	(0.028)	(0.029)	(0.030)	(0.030)	(0.034)
Obs.	424	410	403	374	344
R^2	0.956	0.949	0.942	0.941	0.943

Notes: The table presents horse-race regressions of Greenbook/Tealbook forecasts for core CPI inflation and unemployment for h quarter ahead against the imbalances for FOMC members and staff. We control for the dependent variable lagged two meetings (roughly one quarter) in all regressions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table D.20: Regressions of inflation on imbalances (with expectations)

	All		**FOMC**		**Staff**	
	Core	*Headline*	*Core*	*Headline*	*Core*	*Headline*
I_d	0.067**	0.123***	0.083***	0.142***	0.049*	0.077*
	(0.027)	(0.047)	(0.026)	(0.046)	(0.027)	(0.045)
I_s	0.132***	0.259***	0.094***	0.197***	0.113***	0.236***
	(0.033)	(0.053)	(0.028)	(0.058)	(0.030)	(0.045)
I_e	0.092	0.106	0.136*	0.254**	0.006	-0.170
	(0.108)	(0.148)	(0.076)	(0.127)	(0.094)	(0.154)
π_{t-2}	0.979***		0.981***		0.985***	
	(0.018)		(0.019)		(0.018)	
π_{t-2}		0.960***		0.964***		0.970***
		(0.020)		(0.023)		(0.020)
Obs.	634	634	634	634	634	634
R^2	0.979	0.957	0.978	0.956	0.978	0.956

Notes: The dependent variables are monthly levels of headline and core inflation, expressed as year-over-year percentage changes, from the FRED. Demand, supply, and expectations imbalances are defined as in Equation (14). We compute imbalances separately for FOMC members, staff, and the pooled set of reasons. See Appendix F.3 for more details. Regressions are estimated by OLS with Newey–West standard errors with 8 lags (in parentheses). *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table D.21: Augmented Taylor Rule Estimates (with expectations)

	Full sample	*Pre-Volcker*	*Volcker–GFC*	*Post-GFC*
ϕ_π	1.3***	0.65**	1.76***	3.51*
	(0.34)	(0.27)	(0.21)	(1.87)
ϕ_x	0.99***	1***	0.65***	0.25
	(0.24)	(0.33)	(0.17)	(0.38)
ϕ_s	0.28	1	-0.05	-0.26
	(0.51)	(0.63)	(0.33)	(1.47)
ϕ_d	1.83**	0.21	1.77***	0.81
	(0.9)	(0.75)	(0.53)	(1.41)
ϕ_e	4.35**	4.57**	3.03	-6.17
	(2.01)	(2.13)	(2.4)	(4.13)
ρ	0.87***	0.88***	0.62***	0.92***
	(0.03)	(0.04)	(0.07)	(0.04)
N	634	266	229	139
R^2	0.932	0.919	0.902	0.960

Notes: Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (17), for different samples: *Full sample* = 1960Q1–2025Q2; *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. φ_π is the response to π_t (core inflation), φ_x to x_t (log output gap); φ_d, φ_s, and φ_e to the demand, supply, and expectations textual imbalances. ρ is the interest-rate smoothing parameter. Eq. (18) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table D.22: Interest rates, yields, and term premia's response controlling for macro news (with expectations)

	FFR/ Shadow Rate	*Risk Neutral Yields*		*Term Premia*	
	i_t^s	1y	5y	1y	5y
I_s	0.100**	0.059*	0.029	0.025	0.048***
	(0.047)	(0.035)	(0.034)	(0.015)	(0.018)
I_d	0.184**	0.246***	0.234***	-0.001	0.032
	(0.089)	(0.073)	(0.073)	(0.020)	(0.023)
I_e	-0.065	-0.056	-0.004	0.048	0.010
	(0.142)	(0.124)	(0.109)	(0.038)	(0.049)
Obs.	284	284	284	284	284
R^2	0.983	0.978	0.965	0.883	0.970

Notes: The table re-estimates regressions of the policy rate, risk-neutral yields, and term premia on demand, supply, and expectations imbalances, with and without the set of macro-news predictors of Bauer and Swanson (2023). Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and the level of the dependent variable at $t-2$. See the main text for details on the variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table D.23: High-Frequency Monetary Policy Shock Regressions (with expectations)

	MPS	*S&P 500*	*ED1*	*ED4*	*VIXY*
I_c	0.018*	-0.233***	0.018**	0.014	0.012***
	(0.010)	(0.064)	(0.008)	(0.011)	(0.002)
I_d	0.002	-0.019	0.001	0.001	-0.002
	(0.005)	(0.050)	(0.003)	(0.007)	(0.002)
I_e	0.002	-0.228***	0.000	0.002	0.009***
	(0.015)	(0.085)	(0.011)	(0.017)	(0.003)
Obs.	184	184	184	184	95
R^2	0.186	0.077	0.126	0.178	0.190

Notes: The table reports regressions of high-frequency changes in asset prices and the monetary policy shock of Bauer and Swanson (2023) on the demand, supply, and expectations imbalances extracted from FOMC statements. Asset price changes are from ten minutes before to twenty minutes after the announcement. The dependent variables include the change in the S&P 500 e-mini futures, one- and two-quarter ahead Eurodollar futures (ED1 and ED4), VIXY, and the monetary policy shock from Bauer and Swanson (2023). VIXY is an exchange-traded fund tracking short-term VIX futures; changes in VIXY are computed from TAQ data. The control set corresponds to the macro news predictors in Bauer and Swanson (2023), described in Section 4.6. Sample runs from May, 1999, to December, 2023. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

E Output Examples

This appendix reports the full LLM classification output for four illustrative reasons discussed in the main text: from the meetings on September 16, 2008; June 15, 2022; a late-1960s meeting; and a 1970s meeting during the oil shock.

E.1 Example 1: Day After Lehman's Collapse (2008-09-16)

Passage: Participants "generally thought that the outlook for inflation had improved, mainly reflecting … the weakening of the labor market."

Importance score: 6.0

Speaker: FOMC

All reasons: Y

Reason (label): weakening labor market

Category: Demand

Inflation effect: Down

Output effect: Down

Scratch: weak labor market $\Rightarrow$ cooler demand $\Rightarrow \downarrow \pi, \downarrow y$

E.2 Example 2: Covid (2022-06-15)

Passage: Participants noted that "Most agreed that risks to inflation were skewed to the upside and cited several such risks, including those associated with ongoing supply bottlenecks and rising energy and commodity prices."

Importance score: 7.0

Speaker: FOMC

All reasons: Y

Reason (label): ongoing supply bottlenecks

Category: Supply

Inflation effect: Up

Output effect: Down

Scratch: supply bottlenecks constrain output/costs $\Rightarrow\uparrow \pi, \downarrow y$

E.3 Example 3: Fiscal Stimulus in Late 1969 (1969-12-16)

Passage: Staff emphasized that "Prices are continuing to rise at a rapid rate, and the forces making for further inflation appear impressively strong. The fiscal outlook is in a state of disarray, with all of the questions to be resolved concerned with how much more stimulative fiscal policy will become."

Importance score: 8.0

Speaker: Staff

All reasons: Y

Reason (label): fiscal outlook turning more stimulative

Category: Demand

Inflation effect: Up

Output effect: Up

Scratch: fiscal stimulus $\Rightarrow\uparrow$ demand $\Rightarrow\uparrow \pi, \uparrow y$

E.4 Example 4: Materials Shortages During the Oil Shock (1973-10-16)

Passage: Staff stated that "In this kind of environment, inflationary pressures will surely remain strong. Persisting shortages in the basic materials-producing industries will make possible upward adjustments in prices to reflect increases in costs and to take advantage of any easing in Phase IV price restraints."

Importance score: 8.0

Speaker: Staff

All reasons: Y

Reason (label): basic materials shortages

Category: Supply

Inflation effect: Up

Output effect: Down

Scratch: materials shortages $\Rightarrow\uparrow$ costs $\Rightarrow\uparrow \pi, \downarrow y$

F Robustness and Extensions

This appendix presents one robustness exercise and two extensions. Appendix F.1 examines how the results change under different functional forms for the imbalances. Appendix F.2 compares the assessments of FOMC members and staff. Appendix F.3 adds imbalances based on expectations to the analysis.

F.1 Different Functional Forms

This appendix examines how alternative functional forms for the imbalance measures affect our main results. Given a set of categories $X = \{x_1, x_2, x_3, ..., x_N\}$, we define the vector of net differences

$$N_t = \begin{bmatrix} N^{(+)}_{x_1,t} - N^{(-)}_{x_1,t} \\ N^{(+)}_{x_2,t} - N^{(-)}_{x_2,t} \\ \vdots \\ N^{(+)}_{x_N,t} - N^{(-)}_{x_N,t} \end{bmatrix} \tag{65}$$

The normalization in the main text can be interpreted as the L^1 normalization,

$$I_t = \begin{bmatrix} I_{x_1,t} \\ I_{x_2,t} \\ \vdots \\ I_{x_N,t} \end{bmatrix} = \frac{N_t}{\|N_t\|_1}. \tag{66}$$

We consider two alternative normalizations. The first divides the net imbalances by a moving average of the total number of reasons over the last eight meetings,

$$I^{ma}_t := \frac{N_t}{\sum_{j=0}^{7}(\sum_{i=1}^{N} N^{(+)}_{x_i,t-j} + N^{(-)}_{x_i,t-j})/8}. \tag{67}$$

The second applies the inverse hyperbolic sine to the net imbalances,

$$I_t^{asine} := asine(N_t) = \log\left(N_t + \sqrt{N_t^2 + 1}\right). \tag{68}$$

Appendix Tables D.9, D.10, and D.12 replicate Tables 2, 3, and 7 using the moving-average normalization. Appendix Tables D.14, D.15, and D.16 report the same regressions using the inverse hyperbolic sine normalization.

F.2 Comparing FOMC members and staff

FOMC members do not make decisions in isolation. The Board of Governors' staff act as the FOMC's information and policy engine: they gather real-time data and intelligence, synthesize it with models, and circulate these materials before each meeting. Staff economists then present the economic and financial situation to inform the Committee's policy choice.

The question of where FOMC members versus staff matter has practical implications for monetary policy. Svensson (1999) clearly separates their roles: staff construct inflation and output-gap forecasts under alternative policy scenarios, and the FOMC chooses among these scenarios. Romer and Romer (2008) compare FOMC and staff forecasts and find that FOMC forecasts add little information beyond staff forecasts, while the forecast difference predicts monetary policy shocks. Bennani, Kranz, and Neuenkirch (2018) document substantial disagreement between the policy rate implied by the staff's Greenbook and the rate chosen by the FOMC.

We use our measure to examine whether staff and FOMC perceptions of the causes of inflation are complementary or whether one group's assessment has a disproportionate impact. As described in Appendix A, we flag each inflation reason as presented by either an FOMC member or staff. We then compute imbalances as in Section 2 separately for each group.

Figure C.11 plots the share of reasons discussed by FOMC members over time. The share declines under Burns, then rises to about 80% under Miller and remains high under Volcker and Greenspan. It falls under Bernanke, especially after the GFC, then rises modestly under Yellen and more sharply under Powell.

Table D.17 reports the regressions in Table 2 estimated separately for FOMC members and staff. Coefficients are similar across groups, but the demand coefficient for staff is only significant at the 10% level. Table D.18 reports the specification in Table 7 estimated separately for FOMC and staff. The policy-rate loadings on demand and supply imbal-

ances are smaller when based on staff reasons and statistically significant only for FOMC reasons. Term-premium loadings on supply imbalances display the same pattern: they are larger and statistically significant only for FOMC reasons.

Appendix Table D.4 repeats Taylor rule estimation by monetary policy period separately for imbalances constructed solely from FOMC participants and staff members. We see that the greater responsiveness of the policy rate to perceived inflationary demand imbalances come mostly from the FOMC members. In comparison, the policy rule is much less sensitive to the assessments made by staff members.

Table D.19 reports horse-race regressions of Greenbook forecasts for inflation and unemployment on demand and supply imbalances for both FOMC and staff, controlling for the dependent variable lagged two periods. Inflation forecasts load more on FOMC supply imbalances than on staff supply imbalances. For both groups, loadings on demand imbalances are much smaller, especially for staff. As expected, demand imbalances are strongly associated with lower unemployment forecasts for both groups, whereas supply imbalances are roughly unrelated.

The results suggest that both staff and FOMC assessments are relevant and help explain forecasts and policy decisions. However, policy rates load much more strongly on FOMC assessments, indicating that FOMC members' views carry disproportionate weight in the final decision.

F.3 Including Expectations

Inflation expectations are difficult to measure, and the way FOMC participants discussed them changed substantially over time. We constrain the LLM to identify long-run expectations rather than near-term forecasts and to capture shifts in expectations, not simple statements of anchoring. After generating expectation-related reasons as described in Appendix A, we restrict the baseline specification to cases explicitly referring to long-run expectations and their changes, which we incorporate into the prompt through few-shot learning.

Table D.20 adds the expectations imbalance to the inflation regressions. The coefficients for the expectation term are positive but insignificant in the full sample, and significant only for FOMC members' statements. Table D.21 extends the Taylor rule estimates with the expectations term, which is positive across all subsamples except post-GFC, where the sign reverses but remains insignificant. Table D.22 reports regressions of the shadow rate and term premia controlling for macro news predictors; coefficients on expectations are insignificant. Table D.23 presents high-frequency monetary policy shock

regressions using statement-based imbalances. In this case, expectations imbalances have a negative and significant effect on S&P 500 futures and a positive and significant effect on volatility, alongside supply pressures.

This narrow definition of expectations-driven inflation likely induces type I error. Although the 1960s and 1970s are widely viewed as periods of intense discussion about expectations, Figure 1 shows relatively few such instances. A key challenge is linguistic: the Fed often used nonstandard terms such as "inflation psychology" or "market psychology" rather than explicit measures of expectations (Gorodnichenko, 2021). Prior research documents how FOMC discussions and policymakers' speeches emphasized inflation psychology before the Great Disinflation (Goodfriend and King, 2005; Rotemberg, 2013; Reis, 2022*b*). We therefore refine the prompt to include this historical language and relax the restriction, again using few-shot learning, for the main analysis based on minutes and transcripts. Appendix A.3.2 provides details. Results remain similar under this specification.

F.4 Entity neutering and look-ahead bias

A central appeal of using LLMs in this paper is that they allow us to translate long meeting transcripts into structured, econometrically usable variables. A corresponding concern is look-ahead bias. Modern LLMs are trained on large corpora that include text written long after many of the meetings in our sample. As a result, even when the model is shown only contemporaneous minutes, it may infer the meeting's identity from contextual cues and implicitly lean on ex-post knowledge to classify the passage.

Our baseline pipeline is designed to limit this channel. First, we instruct the model to base its judgments solely on the provided text and to justify its output with explicit textual support. Second, and more importantly, we separate extraction from classification. In the first stage the model reads the full meeting text but performs only a constrained identification task (listing reasons and quoting supporting language). In the second stage the model is given only the specific reason–context pair and is asked to classify it using an "Econ 101" sign-restriction logic. By design, the classifier never sees the full minutes, which reduces its ability to identify the underlying meeting and import information that is not in the local passage.

This appendix implements an additional, more aggressive safeguard based on the entity-neutering approach of Engelberg et al. (2025). The goal is to make each reason–context pair effectively unidentifiable with respect to meeting date (and any other uniquely identifying attributes), while preserving the economic content needed for the second-stage

classification.

The procedure starts from the first-stage outputs: for each identified inflation reason we retain (i) a concise reason label and (ii) the accompanying explanation/context excerpt used to justify that label. This is the object that enters the second-stage classifier in the baseline pipeline.

We prompt an LLM to remove or replace any information that could anchor the passage to a particular meeting date or uniquely identify individuals or episodes. Concretely, the model replaces dates, named people, named events, and other sharp identifiers with generic placeholders (e.g., *year_x*, *person_x*, *event_x*), while being explicitly instructed to preserve the economic structure required for classification (supply vs. demand vs. expectations/monetary policy, and the implied directions for inflation and real activity).

Masking alone does not address memorization: even if explicit identifiers are removed, an LLM might still recognize distinctive phrases from its training data. We therefore add a paraphrasing step that rewrites the reason label and the supporting passage so that the token sequence and sentence structure are substantially altered, while keeping the core content fixed. The full prompts used are shown in Appendix A.6.

After masking and paraphrasing, we run an adversarial identification test. A fresh, "omniscient" LLM is shown the neutered text and asked to guess the most likely meeting date. We treat the text as *identifiable* if the guessed date falls within a ±14 day window around the true meeting date. If the adversary succeeds, we re-run masking and paraphrasing more aggressively and repeat the adversarial test. We allow up to nine iterations, mirroring the iterative, adversarial structure emphasized by Engelberg et al. (2025). If the adversary fails at any iteration, we stop and retain that neutered output for downstream analysis; if the adversary continues to succeed through the iteration cap, we drop the corresponding observation.

Results. Applying this procedure to our full set of first-stage reason–context pairs yields complete neutering within the iteration cap: every reason becomes unidentifiable within nine iterations, with the vast majority neutered after a single pass. We then feed the final neutered reason–context pairs into the same second-stage classifier as in the baseline pipeline, construct meeting-level counts and normalized shares exactly as in the main text, and compare these neutered measures to their raw counterparts.

Two results summarize the diagnostic value of this exercise. First, meeting-level demand and supply counts constructed from the neutered text remain very close to the baseline measures: the raw and neutered demand/supply series (split by inflation-increasing vs. inflation-decreasing reasons) exhibit correlations above 90%. Second, repeating the key

empirical specifications using the neutered variables yields qualitatively unchanged conclusions. Together, these findings indicate that our main results are unlikely to be driven by look-ahead bias from meeting re-identification or memorized phrasing, and that the economic content relevant for classification is preserved by the neutering procedure.