In the Fed's Mind*

Ali Kakhbod
UC Berkeley

Amir Kermani
UC Berkeley
& NBER

Bernardo Maciel
UC Berkeley

This version: November 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—an asymmetry consistent with New Keynesian theory but absent from standard Taylor rules. 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. Our approach offers a scalable and transparent method to extract structured economic insights from central bank communication and sheds new light on how the Fed's perceptions influence its policy stance.

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

*First version: June 2025. We thank Michael Bauer, Emi Nakamura and Jón Steinsson 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.

Despite its theoretical importance, empirically 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 may be decomposed from an econometrician's perspective—not necessarily how central banks themselves viewed the shocks in real time.

In this paper, we introduce a new, text-based methodology to uncover the Federal Reserve's real-time attribution of inflation to supply and demand forces. Our approach is rooted in the "narrative" tradition in macroeconomics, as pioneered by Friedman and Schwartz (1963) and expanded by Romer and Romer (1989, 2023), which seeks to infer policymakers' beliefs and motivations from qualitative records. However, unlike earlier narrative approaches—which required labor-intensive manual reading and interpretation—we leverage recent advances in large language models (LLMs) to extract and classify economic content from text systematically and at scale.

We apply this methodology to the full corpus of Federal Open Market Committee (FOMC) meeting records from 1936 through 2025. These records capture the discussions and rationales behind monetary policy decisions, offering a rare window into how policymakers understood macroeconomic conditions in real time. Our method proceeds in two steps. First, we use LLMs to identify and extract specific statements that cite reasons for inflation. Second, we classify each cited reason as supply-driven and demand-driven using a macroeconomically grounded criterion: whether the factor is expected to move inflation and output in the same or opposite directions. For example, strong consumer spending would increase both inflation and output (demand), whereas an oil price spike would raise inflation and lower output (supply).

In the second stage, the model is not provided with the full text. Instead, it receives only a short label summarizing the reason (e.g., "declining consumer spending") together with an accompanying explanation that explicitly references the text (e.g., "The minutes note that 'consumer spending softened appreciably in recent months,' contributing to economic weakness."). The split keeps the classifier blind to the rest of the document, limiting look-ahead bias and hallucination. Aggregating tags delivers counts of demand, supply, and expectations reasons for each meeting.

We then show our methodology aligns with historical narratives of inflation; for instance, it identifies reasons related to surge in energy prices in the 1970s, at the time of the oil shocks, a surge in disinflationary demand reasons during the Great Recession, followed by a reversal during Covid, when inflationary supply-related reasons spike up. We show that our measure relates to contemporary inflation as theory predicts: inflation responds more to supply imbalances than to demand imbalances.

We use our text measure to document empirical facts. We convert meeting-level reason counts into demand and supply "imbalances"—the net share of inflationary minus disinflationary remarks in each category—and show that both covary positively with inflation, with a larger loading on supply than on demand.

Motivated by New Keynesian theory in a simple model, we augment a Taylor-rule specification with empirical demand and supply factors from our inflation diagnosis. Consistent with the model, we find that interest rates respond more to the demand shock than to the supply shock, even though supply shocks contribute to inflation the most – a result absent in the standard rule.

In addition, we find that policy responses change substantially over time. While in the pre-Volcker period interest rates respond weakly to inflation and strongly to supply, the demand response becomes much larger in the Volcker-Greenspan era and after. This is in line with the standard historical perspective of monetary policy becoming more sophisticated with Fed Chair Paul Volcker.

Consistent with this mechanism, yield-curve decompositions show that demand imbalances mainly raise risk-neutral yields (a tighter expected path of short rates), whereas supply imbalances raise term premia (greater policy-path uncertainty). We also show that the inflation impact of demand imbalances fades faster than that of supply, and that standard pre-announcement macro news explains only a modest share of the meeting-to-meeting variation in the Fed's diagnosis—highlighting that the FOMC records carry independent additional information.

Contribution to the Literature. Our paper contributes to three main strands of the literature. First, recent literature has tried to empirically decompose inflation in demand and supply components. An early seminal paper is Smets and Wouters (2007), which uses a DSGE approach to decompose inflation into productivity, demand, monetary policy shocks, and mark-ups. More recently, Shapiro (2024) uses the sign restriction of demand and supply shocks to identify demand and supply contributions to categories of the personal consumption expenditures (PCEs). Eickmeier and Hofmann (2022) utilize the sign restriction to estimate a factor model, using several macro series to decompose headline inflation into demand and supply components. However, this approach uses information that is not necessarily available to policy makers at the time of their decision and, furthermore, might not be the information they are focusing on. Instead, our methodology focuses on what the FOMC participants are discussing before making their decisions.

Second, our paper relates to the empirical literature on policy rules based on inflation targeting, focusing on the Taylor rule (Taylor, 1993) and its variations. Clarida, Galí, and Gertler (2000) show that the interest rate response changed significantly over time. Orphanides (2001, 2002) show that the Fed's misperception of real-time macroeconomic data might have contributed to poor decisions in the 1970s, even though the policy rule could have been sensible ex-ante.

Even though in its simplest form the Taylor rule implies interest rates respond to inflation irrespective of its perceived causes, standard New Keynesian models imply that the Fed should react more to inflation when it is perceived to be caused by demand and react less if it is perceived as being supply-driven. Some recent theoretical work have explored this type of policy response asymmetry. Beaudry, Carter, and Lahiri (2023) present a model with bounded rationality in which the Fed accomodates supply shocks up to a threshold, after which it becomes hawkish. Guerrieri et al. (2021) show a model in which sectoral changes in demand occurs as an endogenous cost-push shock, which breaks divine coincidence. Nakamura, Riblier, and Steinsson (2025) highlight the importance of central bank credibility in enabling them to look through cost-push shocks without destabilizing inflation expectations. A recent related empirical paper is Hofmann, Manea, and Mojon (2024), which augments a Taylor rule with the demand and supply components of inflation produced by Shapiro (2024). We perform a similar exercise, but with our own measures of the demand and supply contributions to inflation, which are based directly on the Fed's discussions rather than on macroeconomic series.

Third, an alternative strand of the literature examines central bank communication around policy announcements and policy discussions in the FOMC meetings, using natural language processing or narrative analysis to identify the origins of these shocks.

Lucca and Trebbi (2009) use unsupervised semantic-orientation to compute "hawkishness" scores from FOMC statements. Hansen, McMahon, and Prat (2017) use latent Dirichlet allocation (LDA) to measure the forward guidance communicated in FOMC statements. Gorodnichenko, Pham, and Talavera (2023) analyze the tone of FOMC press conference audio and find that markets respond positively to a dovish tone. Similarly, Curti and Kazinnik (2023) analyze FOMC press conference videos to quantify the Chair's facial expressions and find that investors react adversely to negative expressions. Handlan (2022) applies machine learning to the text of FOMC announcements to forecast changes in federal funds rate futures. Hansen and Kazinnik (2023) uses ChatGPT to parse and classify the policy stance of FOMC announcements, comparing it with human classification. Cieslak et al. (2023) exploits FOMC transcripts to document the relationship between FOMC members' perceptions of higher inflation uncertainty and their policy stance. [1]

Our methodology contributes to the monetary policy shocks literature by quantifying a new dimension that is key to the central bank's decision: their perception about the causes of inflation. This can help explain *why* the central bank surprised the market when it moved interest rates. As Bauer, Pflueger, and Sunderam (2024*a*,*b*) point out, the market's perception about the Fed's response is critical for monetary transmission and monetary policy surprises. We show that understanding the Fed's own perception about the causes of inflation is also important in helping us understand the changes in policy responses.

The remainder of the paper is organized as follows. 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 3 presents the empirical results. Section 4 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 each data source, 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.[2] For our empirical tests, however, we focus on the

[1] For other applications of generative AI in finance see Bybee (2023); Fedyk et al. (2023); Kakhbod et al. (2024), among others. For an excellent overview of AI in finance, see also Eisfeldt and Schubert (2024).

[2] 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

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 come from a coherent monetary policy framework.[3]

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).[4]

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

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.

[3]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.

[4]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.

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 quotes to the original text. For each identified reason, the model records three fields: (i) a concise label, (ii) a quotation supporting that label, 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, expectations, monetary policy, or indeterminate/unclear.

More specifically, we provide the model with an objective benchmark, the sign restrictions implied by demand- and supply-side shocks, 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. This is a standard criterion for distinguishing demand and supply shocks and is also present in our theoretical model.

In addition, we classify reasons into two further 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.

In this stage, the LLM's main task is to determine—based on the provided explanation and surrounding context—whether the factor in question moves inflation and output in the same direction (demand) or in opposite directions (supply). Factors 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 additional 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. 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.

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".

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.[5] 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 "China COVID lockdowns worsening supply chains". It refers specifically to a discussion of the lockdowns in China contributing to supply chain disruptions. The full output is also available in Appendix

[5]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).

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".

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 and supply. 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. The common related to supply mention cost-push factors, such as energy prices and productivity.

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 shocks, 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 text-anchored explanation; thus, identical labels can be mapped to different classifications when the surrounding context differs. 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 $N_{d,t}^{(+)}$ and $N_{d,t}^{(-)}$ denote, respectively, the number of demand reasons on date t that the minutes describe as raising or lowering inflation; define $N_{s,t}^{(+)}$ and $N_{s,t}^{(-)}$ analogously for supply. 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.

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

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
Total		763	6.675	4.127	0	4	6	9	35

Notes: For each FOMC meeting ($n = 789$) we count how many demand and supply reasons the minutes 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.

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 counting discussions of output as implicit reasons for inflation; the main conclusions are unchanged.

Figure C.1 plots, for each meeting, the total number of reasons we extract—pooling demand and supply, 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.2 turns to composition. A filled cell appears whenever a demand- or supply-related reason is mentioned, giving a time-series indicator of which categories surface when. We can see that inflationary demand reasons appear throughout history, becoming more absent around the Great Recession; disinflationary demand reasons typically appear during recessions. Indeed, we can see that disinflationary reasons appear more prominently during the Eisenhower Recession (1957 - 1958), the Rolling Adjustment Recession (1960 - 1961), in the early 1980s recession, around the dot-com bust and early 2000s recession, and mainly around the Great Recession and subsequent years.

Inflationary supply reasons are present throughtout the sample, but are relatively more frequent around the 1970s oil shock, the 1990 oil shock, and the energy crises leading up to the Great Recession in the 2000s. Disinflationary supply shocks appear mostly in the

2010s during the decline of energy prices, especially after 2015 (coinciding with the timing of the 2010s oil glut).

2.3 Post-Accord Narratives

Figure 1 shows the annual totals of demand- and supply-related remarks (left axis). For context, the right axis overlays the federal funds rate and inflation. Appendix Figure C.3 shows the net imbalances for demand and supply (increasing minus decreasing counts)

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.

Figure 1: Aggregated Reason Counts Over Time

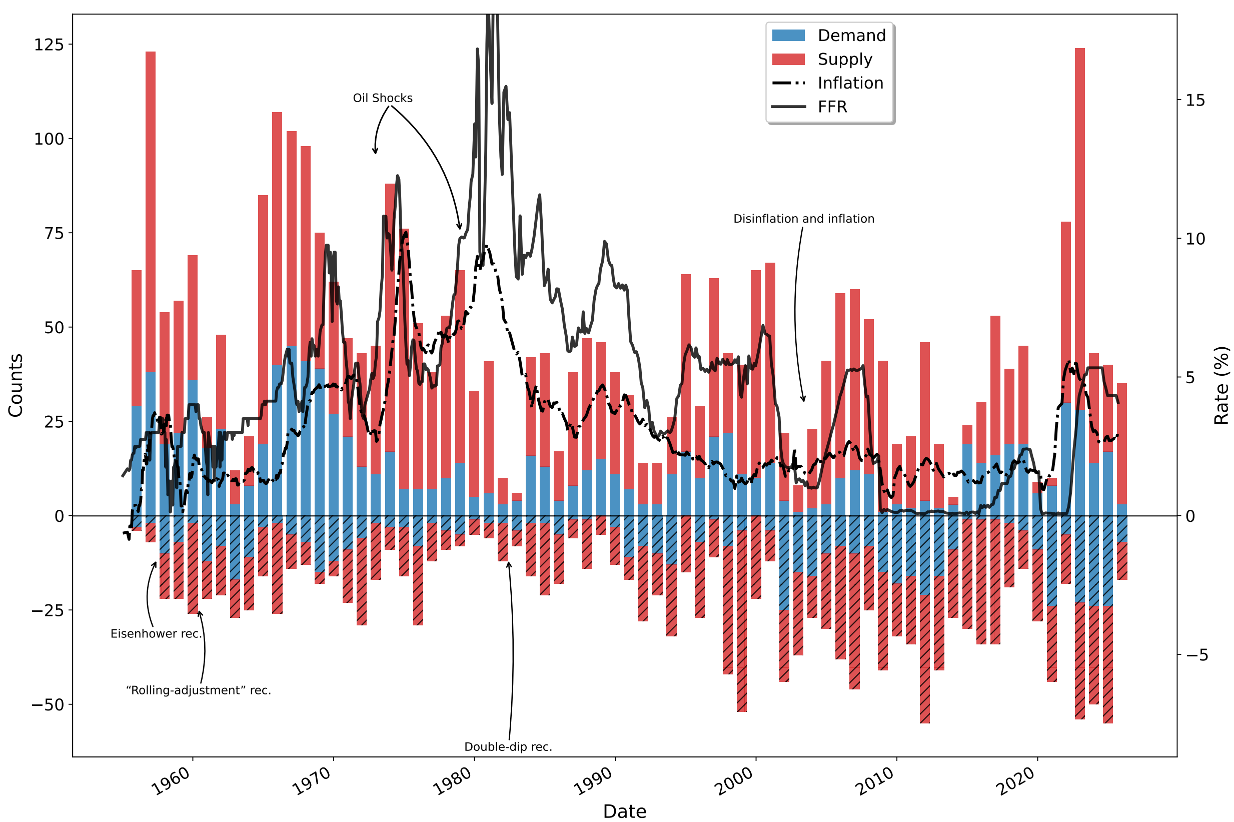

Notes: This figure aggregates the demand and supply reason counts annually, separately for reasons increasing (bars going up, full) and decreasing inflation (bars going down, hatched) in the left axis. Additionally, we present the Federal Funds Rate and inflation in the right axis. We use CPI inflation 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.

Eisenhower recession (1957 – 58). Rate hikes meant to cool earlier price pressures collided with a sudden slump in autos, steel, and housing, worsened by an overseas liquidity squeeze and the 1957 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.

"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 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, which is mostly composed of inflationary supply reasons. 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 supply shocks, 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), 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 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 fed funds rate responds strongly to demand-driven inflation and partially looks through supply shocks, 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 shocks raise risk-neutral yields, whereas supply shocks raise term premia. This finding is consistent with the New Keynesian stabilization trade-off: supply shocks force the central bank to balance inflation against economic activity, increasing uncertainty about the future policy path (and thus term premia), while demand shocks 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.

3.1 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.1 shows, 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

$$I_{x,t} := \frac{N_{x,t}^{(+)} - N_{x,t}^{(-)}}{\sum_{z \in \{d,s\}} \left| N_{z,t}^{(+)} - N_{z,t}^{(-)} \right|}. \tag{1}$$

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 imbalance measure is in 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 all other reasons discussed in that meeting.

3.2 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:

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

where π_t is either core or headline inflation, and the imbalances are constructed separately only for staff, only for FOMC members, and all reasons. As an initial validation, we expect the imbalances to co-move positively with inflation.

Table 2: 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 (1). 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.

3.3 Taylor rule augmented with demand and supply imbalances

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:

$$i_t = \beta_0 + \rho i_{t-2} + (1-\rho)\left[\beta_\pi \pi_t + \varphi_x x_t\right] + \varepsilon_t. \tag{3}$$

Even though explanations for why such a smoothing exists differs in the first place, this is a widely accepted feature of monetary policy in the real world.[6] 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.[7] 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.

Hofmann, Manea, and Mojon (2024) (HMM) augment this specification with demand and supply components of inflation using the decomposition from Shapiro (2024). Similarly, we define a Taylor rule that includes our demand and supply components:

$$i_t = \beta_0 + \rho i_{t-2} + (1-\rho)\left[\beta_\pi \pi_t + \varphi_x x_t + \varphi_d I_{d,t} + \varphi_s I_{s,t}\right] + \varepsilon_t. \tag{4}$$

For this and throughtout the rest of our main analysis, we use all reasons to construct the imbalances, without distinguishing between FOMC participants and staff.

[6] 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.

[7] 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".

Following the methodology of Carvalho, Nechio, and Tristão (2021), we can estimate the structural specification of Equations (3) and (4) by OLS in two steps. First, we define the auxiliary reduced-form equation:

$$i_t = \beta_0^{aux} + \rho i_{t-2} + \varphi_\pi^{aux} \pi_t + \varphi_x^{aux} x_t + \varphi_d^{aux} I_{d,t} + \varphi_s^{aux} I_{s,t} + \varepsilon_t. \tag{5}$$

We then estimate Equation (5) by OLS. Finally, we obtain the structural coefficients as

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

and compute standard errors using the delta method.

3.4 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.[8] 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 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.[9]

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.

The second column covers the pre-Volcker era—Chairs Martin (1951–70), Burns (1970–78), and Miller (1978–79).[10] The baseline coefficient on inflation is below unity, consistent

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

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

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

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 (4), 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. (5) 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.

with the accommodative stance documented by Clarida, Galí, and Gertler (2000). The imbalance coefficients are both small and insignificant, suggesting the absence of a targeted policy response during this period.

The third column restricts the sample to the Volcker–Greenspan era until the GFC and yield 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 (shadow) 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 next, imbalances continue to co-move with the yield curve, linking the Fed's perceptions of inflation's sources to market expectations about the future path of interest rates.

Appendix Table D.4 repeats the same estimation 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.

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 shocks. 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 demand-side inflation while partially looking through supply shocks. After the Great 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 (2) with block-averaged imbalances—the mean of contemporaneous through three-meeting lags (L0–L3) and the mean of four–seven-meeting lags (L4–L7)—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.

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 "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.4 plots the estimated coefficients. Both

Table 4: Response of inflation to lagged imbalances

	Dependent variable: π_t			
	Core		*Headline*	
	(1)	(2)	(3)	(4)
Demand L0–L3	0.146	0.044	0.281*	0.225
	(0.098)	(0.217)	(0.161)	(0.294)
Supply L0–L3	0.469***	0.527***	0.831***	1.059***
	(0.101)	(0.157)	(0.148)	(0.212)
Demand L4–L7		0.095		0.007
		(0.162)		(0.232)
Supply L4–L7		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 ($L0$–$L3$) and, where included, the mean of 4–7 meeting lags ($L4$–$L7$). Columns (1) and (2) use core inflation, while columns (3) and (4) use headline inflation. Columns (1) and (3) include only $L0$–$L3$ 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 $L0$–$L3$ and $L4$–$L7$ 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.

imbalances exhibit similar persistence in our sample; demand and supply are equally as persistent in the Fed's discussions about the causes of inflation.

3.5 Yield curve responses after 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.[11] 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

$$rn_t^{(n)} = y_t^{(n)} - tp_t^{(n)}. \tag{6}$$

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 (5). 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.

Table 5 presents the results using the entire sample, covering August, 1961 to June, 2025. The left panel shows the results for 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. We see a very different picture in the right panel, which shows the effects on term premia. 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.

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. 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.

The asymmetric response to demand and supply imbalances is in line with standard New Keynesian intuition. 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 more uncertain and is reflected in higher term premia, especially at longer maturities, with

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

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)
Obse.	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 (5). 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 $p < 0.01$.

a smaller effect on risk-neutral yields.

3.6 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.[12]

[12] 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

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% 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.

Table 6: Imbalance Regressions on Macro News

	I_d	I_s
NFP	0.079***	0.077***
	(0.022)	(0.025)
S&P500	1.180***	-1.346**
	(0.424)	(0.609)
Slope	-0.022	0.138
	(0.070)	(0.101)
BCOM	0.793*	3.351***
	(0.417)	(0.606)
TR Skew	-0.300**	0.191
	(0.146)	(0.232)
NFP Surp.	0.020	0.039
	(0.024)	(0.025)
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 $p < 0.01$.

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. Estimates are qualitatively unchanged, and in some cases stronger. Demand imbalances are more

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.

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.

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	2y	5y	2y	5y
I_s	0.082*	0.034	0.024	0.032*	0.046***
	(0.044)	(0.032)	(0.030)	(0.018)	(0.017)
I_d	0.150**	0.212***	0.197***	0.008	0.030
	(0.074)	(0.067)	(0.064)	(0.020)	(0.020)
S&P500	-0.465	0.587	0.712*	-0.028	0.072
	(0.539)	(0.440)	(0.418)	(0.188)	(0.180)
BCOM	-0.040	0.161	0.238	0.374**	0.292
	(0.543)	(0.373)	(0.337)	(0.173)	(0.181)
NFP	0.028*	0.011	0.014	0.005	0.002
	(0.016)	(0.015)	(0.016)	(0.005)	(0.006)
NFP Surp.	0.095***	0.075***	0.064***	-0.002	-0.012**
	(0.019)	(0.020)	(0.018)	(0.004)	(0.006)
Slope	-0.618***	-0.477***	-0.351***	0.281***	0.562***
	(0.108)	(0.086)	(0.083)	(0.023)	(0.019)
Tr. Skew	0.207*	0.417***	0.377***	0.090**	0.128***
	(0.116)	(0.116)	(0.112)	(0.036)	(0.031)
π_t	0.089*	0.088**	0.078**	0.002	0.005
	(0.051)	(0.038)	(0.032)	(0.012)	(0.011)
x_t	0.091***	0.087***	0.072***	0.000	0.004
	(0.026)	(0.023)	(0.023)	(0.007)	(0.007)
i_{t-2}^s	0.932***				
	(0.023)				
$rn_{t-2}^{(n)}$		0.905***	0.897***		
		(0.023)	(0.023)		
$tp_{t-2}^{(n)}$				0.940***	0.961***
				(0.020)	(0.012)
Obs.	284	284	284	284	284
R^2	0.983	0.974	0.965	0.929	0.971

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 macro-news predictors of Bauer and Swanson (2023). 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 $p < 0.01$.

In conclusion, the text-based imbalance series closely track monetary history: supply-side imbalances peak in recognized cost-push episodes; demand-side imbalances mirror the business cycle and reverse sign in downturns as policy counters slack. Both imbalances predict subsequent inflation, with a more persistent link for supply. Although past macro variables partially explain these diagnostics, the imbalances retain incremental predictive power.

3.7 Monetary Policy Shocks

Next, we estimate how imbalances map into asset-price movements around FOMC announcements. Table 8 reports OLS regressions of three-day changes in risk-neutral yields and term premia on demand and supply imbalances, with and without the macro-news predictors of Bauer and Swanson (2023). We also report analogous regressions on the monetary policy shocks of Romer and Romer (2004) updated by Acosta (2023) on the same imbalances, again with and without these controls.[13]

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)
S&P500		0.073		-0.010		-0.007		0.115		-0.225*
		(0.083)		(0.118)		(0.039)		(0.110)		(0.133)
BCOM		0.107*		0.096		-0.027		-0.043		-0.060
		(0.062)		(0.078)		(0.044)		(0.071)		(0.172)
NFP		0.001		0.003		-0.001		-0.002		0.002
		(0.002)		(0.003)		(0.001)		(0.003)		(0.010)
NFP Surp.		0.003*		0.005		-0.002		-0.007**		0.040
		(0.002)		(0.003)		(0.002)		(0.003)		(0.049)
Slope		-0.020**		-0.018		-0.003		0.004		-0.099***
		(0.010)		(0.012)		(0.005)		(0.010)		(0.027)
Tr. Skew		0.035***		0.042***		0.002		0.015		0.200***
		(0.010)		(0.013)		(0.010)		(0.021)		(0.043)
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

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 $p < 0.01$.

The earlier asymmetry persists: changes in term premia load positively on supply imbalances, while changes in risk-neutral yields load positively on demand imbalances, though the latter is not significant at longer maturities. With controls, the risk-neutral yield change also loads on supply at the 10% level, and the term-premium change loads negatively on demand at the 10% level. For Romer–Romer shocks, imbalances are posi-

[13] Data from Miguel Acosta's website.

tively related to the shock; supply is significant at 5%. With macro-news controls, both coefficients remain positive, with supply significant at 10% and demand at 5%.

3.8 Comparison to Shapiro (2024)

Recent work by Shapiro (2024) attempts to isolate the demand- and supply-side drivers of inflation from aggregate macroeconomic variables. We benchmark our FOMC text measure against the series in Shapiro (2024), which classifies each item in personal consumption expenditures (PCE) as demand, supply, or "ambiguous" by using sign restrictions.[14]

Shapiro produces three time series: dem_t, sup_t and ambig_t.[15] In addition, the data contains two separate versions of the series: one using headline inflation and another using core inflation. We consider the two variants separately below.

To make the estimates from Shapiro (2024) comparable to ours, we transform our demand and supply imbalances as

$$P_{d,t} = \alpha_d P_{d,t-1} + (1-\alpha_d) I_{d,t} \tag{7}$$

$$P_{s,t} = \alpha_d P_{s,t-1} + (1-\alpha_s) I_{s,t} \tag{8}$$

where α_d and α_s are calibrated to match the persistence of the correspondent PCE-based decomposition. Table 9 presents the regressions of the PCE measures from Shapiro (2024) on our smoothed shares.

[14] The data is available in the author's website: https://sites.google.com/site/adamshap/data-and-code.

[15] The procedure can be summarized as follows. For each PCE category i and month t run rolling regressions of quantity on lags of quantity and price, and of price on lags of price and quantity; let the one–step residuals be $\hat{u}^Q_{i,t}$ and $\hat{u}^P_{i,t}$. Classify i as "demand" if $\text{sign}(\hat{u}^Q_{i,t}) = \text{sign}(\hat{u}^P_{i,t})$, "supply" if the signs differ, and "ambiguous" if either residual is near zero. Let the set of demand and supply categories be denoted by D_t and S_t. Then demand and supply contributions are the expenditure–share–weighted averages of category inflation over the corresponding sets: $\pi^D_t = \sum_{i\in D_t} \omega_{i,t-1}\pi_{i,t}$ and $\pi^S_t = \sum_{i\in S_t} \omega_{i,t-1}\pi_{i,t}$ (analogously for ambiguous).

Table 9: PCE- and text-based decompositions

Panel A: Headline Inflation

	dem_t	dem_t	dem_t	sup_t	sup_t	sup_t
$P_{d,t}$	0.765**		0.584**	0.758		0.394
	(0.322)		(0.255)	(0.537)		(0.382)
$P_{s,t}$		1.974***	1.929***		3.887***	3.856***
		(0.305)	(0.297)		(0.487)	(0.491)
Obs.	481	481	481	481	481	481
R^2	0.052	0.418	0.448	0.016	0.510	0.514

Panel B: Core Inflation

	dem_t	dem_t	dem_t	sup_t	sup_t	sup_t
$P_{d,t}$	0.967***		0.835***	1.107**		0.833**
	(0.294)		(0.219)	(0.463)		(0.358)
$P_{s,t}$		1.933***	1.882***		3.954***	3.903***
		(0.254)	(0.248)		(0.518)	(0.523)
Obs.	481	481	481	481	481	481
R^2	0.089	0.406	0.472	0.036	0.522	0.543

Notes: The table regresses the the PCE-based decomposition of inflation (sup_t and dem_t) from Shapiro (2024) and our text-based measures ($P_{d,t}$ and $P_{s,t}$). The text-based demand and supply measures are transformed as contemporaneous inflation multiplied by the exponentially weighted demand and supply imbalances, $P_{d,t}$ and $P_{s,t}$, to make both sets comparable. Panel A uses headline inflation for all measures and Panel B use core inflation for all measures. Newey-Wests standard-errors are in parentheses (8 lags). The dependent variable is equal to the the shadow rate whenever that rate is available (1990-2022) and the federal funds rate otherwise. *, **, and *** denote significance at the 10%, 5%, and 1%.

Panels A (headline) and B (core) tell the same story. The PCE-based demand (dem_t) and supply (sup_t) series load positively and significantly on their text-based counterparts, with large R^2. The mapping is sharper for supply: sup_t loads primarily on $P_{s,t}$. By contrast, dem_t loads on both $P_{d,t}$ and $P_{s,t}$. The close correspondence between the FOMC-derived and PCE-based decompositions indicates that our imbalances capture economically meaningful distinctions between demand- and supply-driven inflation. However, a key advantage of the text-based approach is that it reflects what policymakers were actually discussing at the time of each meeting, providing a direct link between perceptions and policy decisions.

4 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-driven 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 "looking through" adverse supply shocks.

We document a systematic asymmetry between demand and supply assessments and yield-curve movements. Higher perceived demand pressure predicts an increase in risk-neutral yields; higher perceived supply pressure predicts an increase in term 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.

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 framework 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 sharpen the empirical lens on the inflation process 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.

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*a*. "Changing Perceptions and Post-Pandemic Monetary Policy." Jackson Hole Economic Policy Symposium.

Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam. 2024*b*. "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.

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.

Bybee, J Leland. 2023. "The Ghost in the Machine: Generating Beliefs with Large Language Models." Yale University Working Paper.

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

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.

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.

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*.

Eisfeldt, Andrea L and Gregor Schubert. 2024. "AI and Finance." National Bureau of Economic Research Working Paper 33076.

Fedyk, Anastassia, Ali Kakhbod, Peiyao Li, and Ulrike Malmendier. 2023. "AI and Perception Biases in Investments: An Experimental Study." *UC Berkeley Working Paper*.

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.

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*.

Handlan, Amy. 2022. "Text Shocks and Monetary Surprises: Text Analysis of FOMC Statements with Machine Learning." Working paper, Department of Economics, Brown University.

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.

Kakhbod, Ali, Leonid Kogan, Peiyao Li, and Dimitris Papanikolaou. 2024. "Measuring Creative Destruction." *MIT Sloan Working Paper*.

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

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

Orphanides, Athanasios. 2001. "Monetary Policy Rules Based on Real-Time Data." *American Economic Review*, 91(4): 964–985.

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

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.

Primiceri, Giorgio E. 2006. "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U. S. Postwar Stabilization Policy." *Quarterly Journal of Economics*, 121(3): 867–901.

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

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

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.

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*.

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

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.

Internet appendix

- Section A contains details of the minutes prompt designs for our LLM methodology.
- Section B.1 contains the details of the model.
- Section C contains supplementary figures.
- Section D contains supplementary tables.
- Section E contains examples of the LLM output.

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.[16] 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
```

[16] 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. 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.
```

B A Simple 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). In what follows, we distinguish true realizations from the central bank's perceptions. Since our focus will be on the perceived values, true values carry a star (e.g. x_t^*, π_t^*), while perceived values appear without a star (e.g. x_t, π_t).

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 factors. 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:

$$d_t^* = \rho_d d_{t-1}^* + \varepsilon_{d,t}^* \tag{9}$$

$$s_t^* = \rho_s s_{t-1}^* + a_s \varepsilon_{s,t}^* \tag{10}$$

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.

We assume the central bank has imperfect information about the demand and supply factors. The central bank makes contemporaneous assessments of the factors, d_t and s_t, based on noisy signals, $\varepsilon_{d,t}$ and $\varepsilon_{s,t}$. For parsimony, we assume these signals are unbiased (so $\mathbb{E}_t(\varepsilon_{d,t}) = \varepsilon^*_{d,t}$ and $\mathbb{E}_t(\varepsilon_{s,t}) = \varepsilon^*_{s,t}$). Additionally, we assume d^*_{t-1} and s^*_{t-1} are known at t. Taken together, these assumptions imply that the best estimators for $\mathbb{E}_t(d^*_t)$ and $\mathbb{E}_t(s^*_t)$ are given by

$$d_t = \rho_d d^*_{t-1} + \varepsilon_{d,t} \tag{11}$$
$$s_t = \rho_s s^*_{t-1} + \varepsilon_{s,t} \tag{12}$$

Moreover, the central bank forms model-consistent nowcasts x_t and π_t by substituting the perceived shocks d_t and s_t into the IS and the Phillips Curve:

$$x_t = -\psi(i_t - \mathbb{E}_t\pi_{t+1}) + \mathbb{E}_t x_{t+1} + \kappa_d d_t \tag{13}$$
$$\pi_t = \lambda x_t + \beta\mathbb{E}_t\pi_{t+1} + \phi_s s_t. \tag{14}$$

The policy instrument always is observed ($i_t = i^*_t$), and all forecasts are common knowledge—there is no disagreement between the monetary authority and market participants.

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

$$i_t = \mathbb{E}_t\pi_{t+1} + \Upsilon_s s_t + \Upsilon_d d_t \tag{15}$$

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:

$$0 < \Upsilon_s < \Upsilon_d \tag{16}$$

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:

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

The Fed plugs in its best prediction for d_t^* and s_t^* (i.e., d_t and s_t) into the objective function and treats it as the truth (i.e., it does not incorporate the uncertainty about them into the objective function or actions in any way). This is the typical modeling approach in most of the macroeconomics learning literature (Primiceri, 2006). Since the central bank does not observe innovations perfectly, it also does not observe x_t or π_t perfectly. All of its decisions are based on its best guess about these values, which can be wrong.

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:

$$\mathcal{L} = -\frac{1}{2}\left[\alpha x_t^2+\pi_t^2\right]+\gamma(\pi_t-\lambda x_t-f_t)+F_t \tag{17}$$

where

$$f_t = \beta \mathbb{E}_t \pi_{t+1} + \phi_s s_t \tag{18}$$

$$F_t = -\frac{1}{2}\mathbb{E}_t \left[\sum_{j=1}^{\infty} \beta^j (\alpha x_{t+j}^2 + \pi_{t+j}^2) \right] \tag{19}$$

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

$$x_t = -\frac{\lambda}{\alpha}\pi_t \tag{20}$$

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

$$\pi_t = -\frac{\lambda^2}{\alpha}\pi_t + \beta \mathbb{E}_t \pi_{t+1} + \phi_s s_t \tag{21}$$

$$\implies \pi_t = \frac{\alpha}{\alpha + \lambda^2}\beta \mathbb{E}_t \pi_{t+1} + \frac{\alpha}{\alpha + \lambda^2}\phi_s s_t \tag{22}$$

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:

$$\pi_t = \frac{\alpha}{\alpha + \lambda^2}\beta \mathbb{E}_t \pi_{t+1} + \frac{\alpha}{\alpha + \lambda^2}\phi_s s_t \tag{23}$$

Imposing rational expectations and the usual transversality conditions (which rule out bubbles), we can iterate forward, conditional on the current beliefs about the shocks:

$$\begin{aligned}
\pi_t =& \frac{\alpha}{\alpha + \lambda^2}(\phi_s s_t) + \left(\frac{\alpha}{\alpha + \lambda^2}\right)^2 \beta(\phi_s \rho_s s_t) \\
&+ \left[\frac{\alpha}{\alpha + \lambda^2}\right]^3 \beta^2 (\phi_s \rho_s^2 s_t) + \ldots \\
=& \frac{\alpha}{\alpha + \lambda^2}\phi_s \left[s_t + \frac{\beta\alpha\rho_s}{\alpha + \lambda^2} s_t + \left(\frac{\beta\alpha\rho_s}{\alpha + \lambda^2}\right)^2 s_t + \ldots \right] \\
=& \frac{\alpha\phi_s}{\lambda^2 + \alpha(1 - \beta\rho_s)} s_t
\end{aligned} \tag{24}$$

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

$$\mathbb{E}_t \pi_{t+h} = \frac{\alpha \phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} \rho_s^h s_t, \quad \forall h \geq 0 \tag{25}$$

$$x_t = -\frac{\lambda \phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} s_t \tag{26}$$

$$\mathbb{E}_t x_{t+h} = -\frac{\lambda \phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} \rho_s^h s_t, \quad \forall h \geq 0 \tag{27}$$

where the equations above follow from the FOC and the DGPs from Equations (11) and (12).

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:

$$i_t = \mathbb{E}_t \pi_{t+1} + \underbrace{\frac{1}{\psi} \left[\frac{\lambda(1-\rho_s)\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} \right]}_{\Upsilon_s} s_t + \underbrace{\frac{1}{\psi} \kappa_d}_{\Upsilon_d} d_t \tag{28}$$

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:

$$\lambda^2 + \alpha(1-\beta\rho) - \lambda(1-\rho) > 0 \tag{29}$$

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 Woodford (2009). Thus, under usual assumptions:

$$0 < \Upsilon_s < \Upsilon_d \tag{30}$$

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.

C Supplementary Figures

Figure C.1: Total Reason Counts Over Time

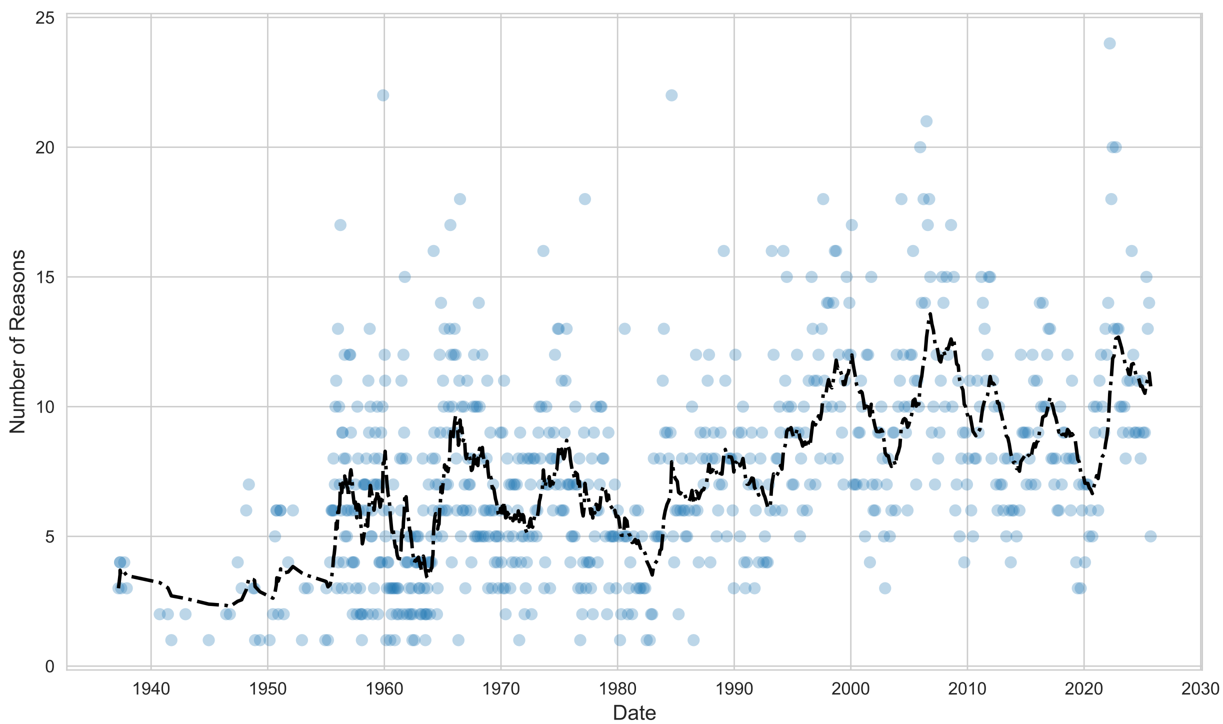

Note: Each dot gives the meeting-level count of reasons—demand or supply, inflationary or disinflationary. The black curve is an exponential moving average with a half-life of eight meetings. See the main text for details on the construction of these variables.

Figure C.2: Presence of Demand and Supply Reasons over Time

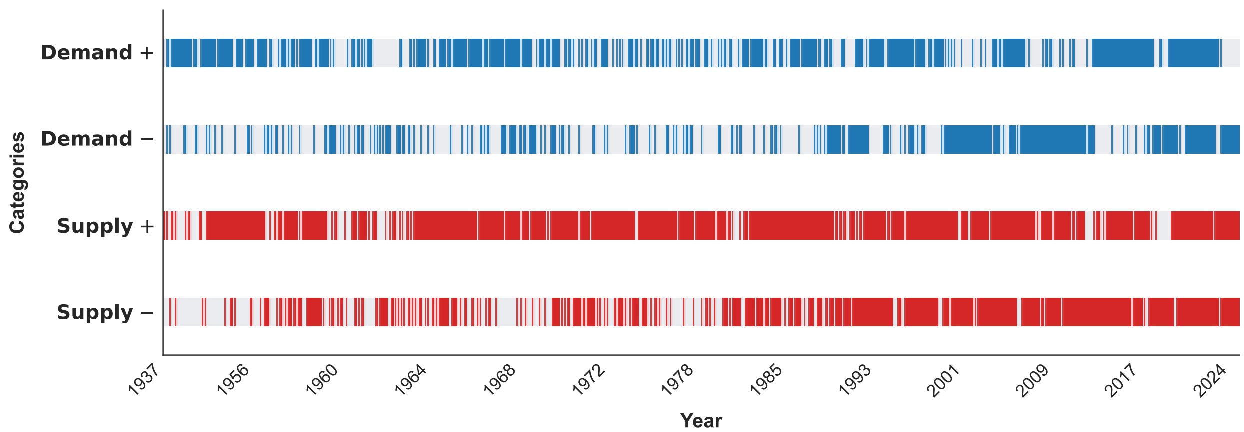

Note: A filled cell marks any FOMC meeting which mentions at least one reason of the indicated category (demanda and supply). Inflationary (disinflationary) reasons are accompanied by a "+" ("−"). See the main text for details on the construction of the variables.

Figure C.3: Net Aggregated Reason Counts Over Time

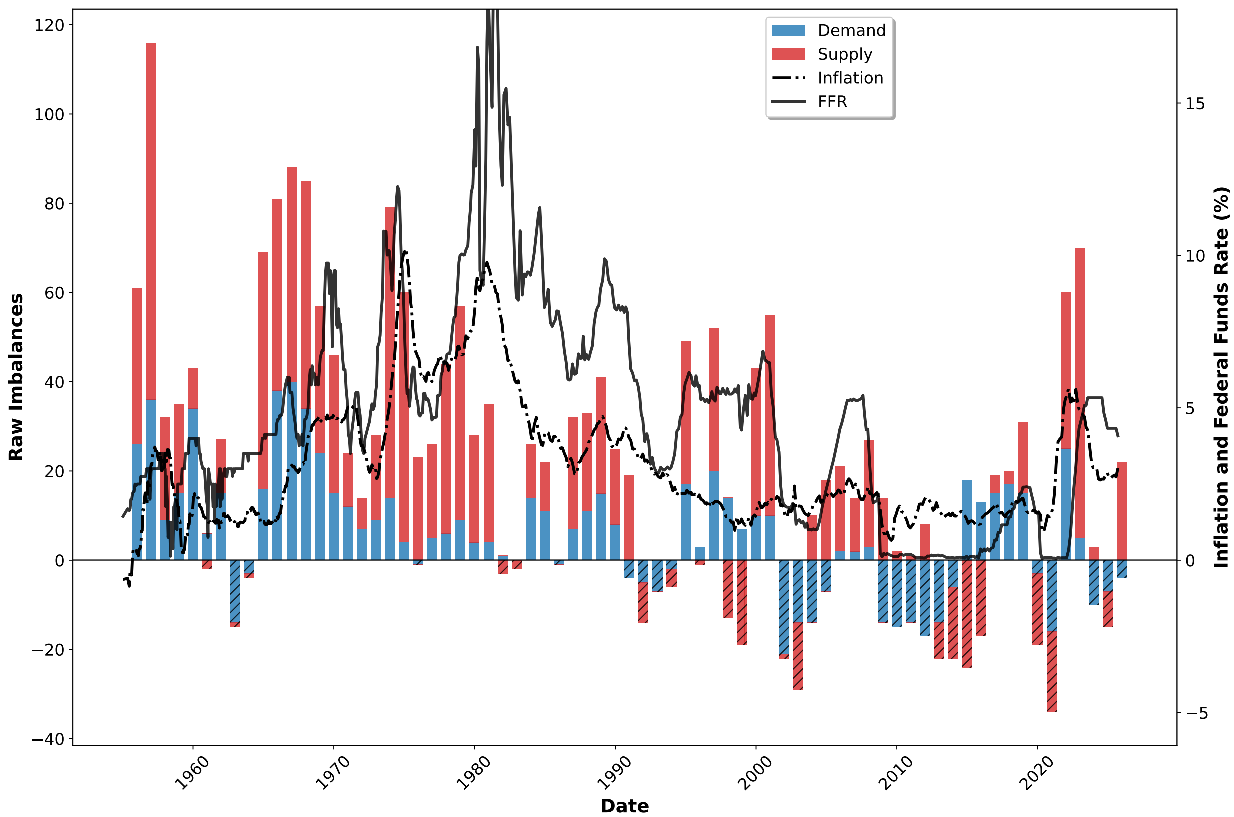

Notes: This figure aggregates the net demand and net supply reason counts annually in the left axis. Net demand (supply) is the sum of demand (supply) reasons contributing to an increase in inflation minus the corresponding reasons contributing to a decrease in inflation in a given year. Additionally, 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.

Figure C.4: Persistence of imbalances

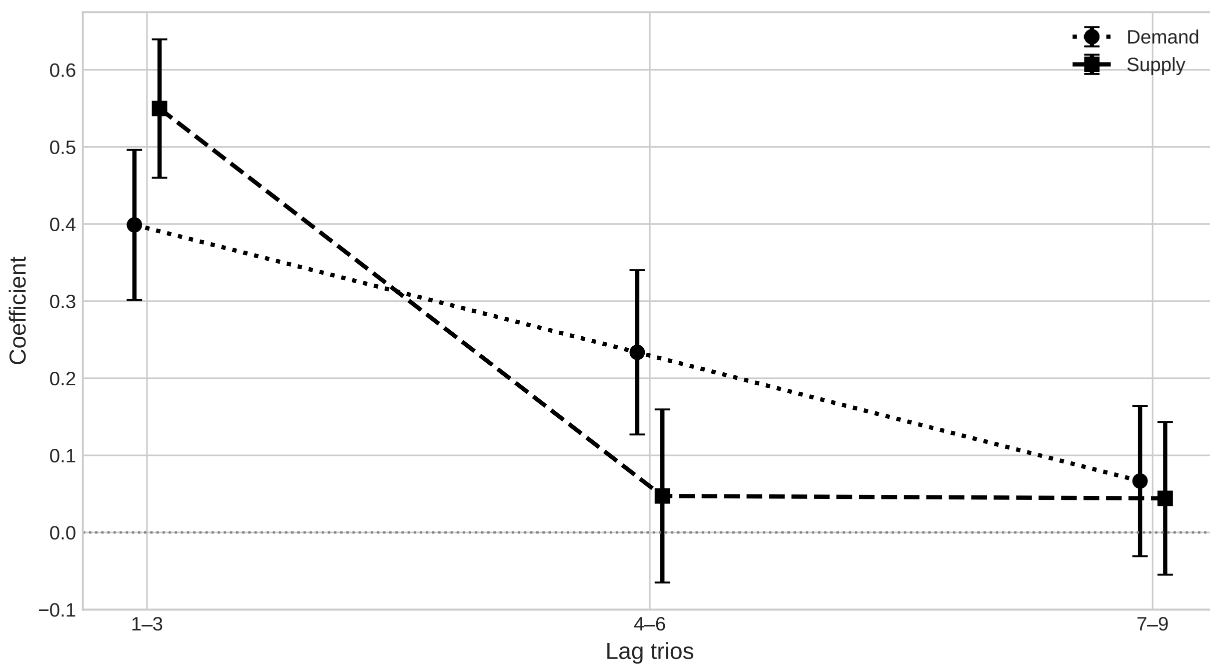

Notes: This figure plots autoregressive coefficients with the associated 95% confidence intervals. Each imbalance is related to the average of adjacent lag trios, estimated separately for demand and supply. Specifically, we regress contemporaneous imbalance on theree subsequent averages: averages of 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.

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
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

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*	
Reason	*Count*	*Reason*	*Count*
resource slack	17	energy prices	66
high resource utilization	16	dollar depreciation	40
slack in labor and product markets	14	productivity gains	26
slack in resource utilization	14	wage increases	21
economic slack	9	food prices	19
tight labor markets	8	oil prices	19
labor market strengthening	7	unit labor costs	18
tightening resource utilization	7	commodity prices	16
underutilized resources (slack)	7	energy price declines	13
fiscal stimulus	6	higher import prices	13
Total	105	*Total*	251

Notes: The table reports the ten most cited explanations for inflation in the minutes, sorted by category. "Demand" and "Supply" 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 (5). 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 (4), 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. (5) 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.

E Output Examples

This appendix contains the full, raw output of the LLM classification for two illustrative reasons discussed in the main text. The first reason comes from the September 16th, 2008 meeting. The second comes form the June 15th, 2022 meeting.

Output Example 1: September 2008

```
Date:          2008-09-16
speaker:       FOMC
category:      Demand
importance_score:   6.0
all_reasons:   Y
reason:        weakening labor market
explanation:   Participants 'generally thought that the outlook for inflation had
improved, mainly reflecting ... the weakening of the labor market.'
scratch:       weak labor market ⇒ cooler demand ⇒ ↓π, ↓y
output_effec:  Down
inflation_effect:   Down
```

Output Example 2: June 2022

```
Date:          2022-06-15
speaker:       FOMC
category:      Supply
importance_score:   7.0
all_reasons:   Y
reason:        China COVID lockdowns worsening supply chains
explanation:   Participants judged that 'the invasion and related events were
creating additional upward pressure on inflation and were weighing on global
economic activity. In addition, participants indicated that COVID-related
lockdowns in China were likely to exacerbate supply chain disruptions.'
scratch:       China lockdowns ⇒ supply chain disruptions ⇒ ↑π, ↓y
output_effec:  Down
inflation_effect:   Up
```