Deciphering the Roots of Inflation with Minutes Text*

Ali Kakhbod UC Berkeley

Amir Kermani UC Berkeley & NBER

Bernardo Maciel UC Berkeley

June 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 2023 in two steps. First, we extract statements citing causes of inflation. Second, we classify each as supply- or demand-driven using macroeconomic sign restrictions. This produces a high-frequency panel tracking the Fed's evolving diagnosis of inflation.

We find that inflation responds more strongly to supply shocks, 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. 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

*We thank Emi Nakamura and Jón Steinsson for their 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 minutes from 1936 through 2023. These minutes record 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 in the minutes that cite reasons for inflation. Second, we classify each cited reason as supply-driven, demand-driven, or ambiguous, 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). This second stage is not exposed to the

full minutes, only to a short label and explanation detailing each reason. The split keeps the classifier blind to the rest of the document, limiting look-ahead bias and hallucination. Aggregating tags delivers counts of demand and supply 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 deflationary 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 later show that our empirical demand and supply factors derived from our measure are related to the Fed's forecasts about inflation and unemployment. Inflation forecasts load positively on both factors, while unemployment forecasts load positively on supply but negatively on demand. Together, these facts show that our classification can help capture the Committee's real-time view of inflation's roots.

Next, we present a simple model that embeds two latent factors into the Clarida, Galí, and Gertler (1999) framework (CGG). We define a demand factor (d_t) that positively affects the output gap and inflation directly, and a supply factor (s_t) that does not change the output gap but positively affects inflation. We further assume that the Fed does not perfectly observe these shocks, but receives a signal about them; the Fed relies on its best guess about the underlying causes of inflation to implement the optimal interest rate response. We show that, under this scenario, real interest rates will be a linear combination of the deeper demand and supply factors, and will respond more to demand than supply.

Motivated by the model, we augment a simple Taylor rule equation with empirical demand and supply factors derived from our measure. As the model indicates, 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 and indistinctly to both shocks, the demand response becomes much larger in the Volcker-Greenspan era. This is in line with the standard historical perspective of monetary policy becoming more sophisticated with Fed Chair Paul Volcker.

Our factors help open the black box of high-frequency policy shocks. The standard identification approach in the literature backs out a monetary policy shock from the jump in asset prices within minutes of an FOMC release. We instead ask what the Committee had in mind *before* it moved the rate. Linking the surprises in Treasury yields, stocks, and Fed funds futures to the demand and supply contributions shows that part of the unexpected change reflects the Fed's own diagnosis of demand and supply pressures recorded in the

minutes, not pure noise. The text, not just the tape, helps explain why markets moved.

Contribution to the Literature Our paper contributes to four 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 led to poor decisions in the 1970s, even though the policy rule might 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 recent 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, and then becomes hawkish after that. Guerrieri et al. (2021) show a model in which sectoral changes in demand occurs as an endogenous cost-push shock, which breaks divine coincidence. 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, our paper contributes to the literature on monetary policy shocks. A large literature in monetary economics constructs policy surprises from changes in asset prices in a narrow window around FOMC announcements. Kuttner (2001) uses a scaled version of

the current-month federal funds futures (FF1) as his monetary policy shock. Nakamura and Steinsson (2018) and Gürkaynak, Sack, and Swanson (2005) use scaled versions of FF1 and next month federal funds futures (FF2), along with Eurodollar futures contracts expiring in two to four quarters (ED2, ED3, and ED4). Swanson (2021) includes S&P 500 futures, FF1, FF2, ED1 to ED4, and Treasury yields with maturities of 2, 5, 10, and 30 years, among other series. Brennan et al. (2024) present a broad overview of the methods and the data sources used in this literature. While these studies vary in their choice of series and in how they aggregate them into factors, they all infer monetary shocks from market responses rather than directly identifying the shocks' origins.

Fourth, an alternative strand of the literature examines central bank communication around policy announcements, 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.[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 believe 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 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 2 presents a simple New Keynesian framework with latent demand and supply factors and characterizes the optimal interest-rate response. In

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

Section 3, we construct measures of demand- and supply-side contributions to inflation and embed them in an augmented Taylor rule. Section 4 reports additional analyses and robustness checks. Section 5 calibrates the model to generate impulse-response functions. Section 6 links our textual measure to high-frequency monetary policy shocks. Section 7 concludes.

1 Data

Our empirical analysis combines textual, macroeconomic, and high-frequency financial data from different sources, described in this section. Our baseline dataset is the full corpus of FOMC meeting minutes, spanning 777 meetings from January 1937 to December 2023.[2]

To initialize the exponential smoothing filter that extracts perceived demand- and supply-side inflation contributions, we feed the entire series into the algorithm. For our empirical tests, however, we focus on the modern policy regime that begins after the Treasury–Fed Accord of 1951—specifically, meetings from 1960 onward to align with the PCE core and headline inflation series. Using the full history to smooth the series lets us avoid arbitrary starting-value effects, while restricting 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. Both start in January 1960 and end in Decembe 2023, totaling 624 observations. The main series of interest rates is the monthly federal funds effective rate from the FRED. This series starts in September 1954 and ends in December 2023, totaling 711 observations.

In order 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] We use the series for the period February 1990 through January 2022. After March 2022, the Atlanta Fed—our data source—halted further updates.

We draw on the Federal Reserve Tealbook (formerly the Greenbook) to obtain headline-

[2]The format of the minutes changes over time. Between 1936 and 1967, we use the "Historical Minutes" from the Fed website; the "Minutes of Actions" between 1967 and 1992; and the modern "Minutes" from 1992 onwards.

[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.

CPI and unemployment forecasts from the current quarter (F0) through four quarters ahead (F4). After merging these projections to the FOMC calendar, the inflation series spans November 1979 to December 2019, giving 293 observations. The unemployment forecasts begin earlier—between May 1967 and November 1969—and run through December 2019, yielding 375 to 443 observations across horizons. All series are obtained from the Philadelphia Fed.

Finally, we incorporate the high-frequency monetary policy shock dataset of Bauer and Swanson (2023*b*), obtained from the San Francisco Fed. Each series measures the change in an asset's price from ten minutes before to twenty minutes after a policy announcement. Our sample includes: (i) federal-funds futures one and two months ahead, available from January 2010 to December 2023 (112 observations); (ii) Eurodollar futures one to four quarters ahead and the 2-, 5-, 10-, and 30-year Treasury yields, spanning February 1994 to December 2023 (231 observations); (iii) the "big" S&P 500 contract, from February 1994 to December 2019 (200 observations); and (iv) the "mini" S&P 500 contract, from September 1997 to December 2024 (206 observations). The baseline monetary-policy shock of Bauer and Swanson (2023*b*) and its orthogonalized counterpart run from December 1990 to December 2023, with 232 and 225 observations, respectively. VIXY futures are sourced directly from the TAQ database. Additional details are provided in Section 7.

2 Methodology

We build a measure of the Fed's own diagnosis of inflation by coding every FOMC minute from 1936 to 2023. 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 OpenAI's 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 any economist can merge, rerun, and extend quickly and conveniently. We use OpenAI's 4o-mini version for the first stage and o4-mini for the second stage.

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''). The full prompt is 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 demand-driven, supply-driven, or indeterminate/unclear.

To do so, 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 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.

Thus, the LLM's task in this stage is to infer - based on the explanation provided and the context - whether the factor at hand drives inflation and output in the same direction (demand) or in opposite directions (supply). Finally, reasons that are ambiguous or tautological ("prices increased", without reference to deeper structural causes) are classified as "other". Manual review corrects the few remaining misclassifications. The full prompt is available in the 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 an "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 high-frequency panel of demand, supply, and residual reasons—each tagged as inflationary or disinflationary—that tracks the Fed's stated view of inflation's causes over almost nine decades. There are a total of 3,171 reasons identified, with 1,867 being demand-related, 919 supply-related, and 385 identified as either ambiguous or tautological and tagged as "other", the residual category. The output of the LLM conveniently provides structured data to the econometrician that can be read as a table for further analysis.

To illustrate, we consider the discussion from the 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 seven reasons discussed by the participants in that ocasion. Six of these are negative-demand reasons. The last is a negative-supply reason associated with the decline in prices of energy and other commodities.

Consider one reason for that date, labeled "declining consumer spending". The explanation provided by the LLM is "*The minutes mention that 'consumer spending softened appreciably in recent months', contributing to economic weakness.*" The explanation identifies an exact quote from the minute, which allows for transparency and reproducibility, and avoids hallucination.

The output from this example, available in Appendix D, shows how we allow the LLM to record its own chain-of-thought. This is done in the "scratch" field, which records the model's internal logic. This field serves a twofold purpose: it improves the quality of the classification and allows researchers to verify its internal consistency. In this case, this field notes that weaker consumption cuts both output and prices, so inflation and the output gap are each marked "Down." This chain of thought, in turn, is consistent with the model's conclusion of classifying the reason as "Demand".

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 we highlight is identified as "COVID-related lockdowns". It refers specifically to a discussion of the lockdowns in China contributing to supply chain disruptions. The full output is available in Appendix D.

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".

After detailing the structure of the raw output of our methodology, we now step back to examine the distribution of reason classifications. Table 1 contains the most common reasons identified for demand, supply, and the residual category. The table shows that some of the common reasons classified as demand are related to employment, wage conditions, and consumer spending. The most common reasons related to supply mention energy prices, changes in industrial commodity prices, and supply chain disruptions. The residual bucket captures ambiguous notions –"inflation expectations", "transitory factors"–and tautologies such as "rapid price increases", "rising prices", and "inflationary pressures".

It is worth noting that the "wage increases" reason appears under both demand and supply. This is because our prompts direct the LLM to distinguish between demand- and

Table 1: Top reasons for inflation by category

Demand		Supply		Other	
Reason	**Count**	**Reason**	**Count**	**Reason**	**Count**
high unemployment	53	energy prices	32	inflationary pressures	16
balance of payments deficit	38	higher energy prices	15	low inflation expectations	13
tight credit conditions	38	energy price increases	14	upward pressures on prices	12
tight labor market	29	wage increases	14	rapid price increases	10
inventory accumulation	27	dollar depreciation	11	consumer price increases	7
wage increases	27	labor shortages	10	inflationary expectations	6
tight labor markets	24	declines in energy prices	9	inflation expectations	6
credit expansion	19	energy price declines	9	rising prices	6
increased unemployment	18	supply chain disruptions	9	transitory factors	5
consumer spending	18	industrial commodity prices	8	supply and demand imbalances	5
Total	**291**	**Total**	**131**	**Total**	**86**

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; "Other'' collects ambiguous or tautological references. "Count'' is the number of occurrences across the full sample. See the main text for details on the construction of these variables.

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.

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 deflationary 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, and do the same for the residual category. These raw counts summarize the Committee's perceived balance of inflationary and deflationary perspectives for each category and in each meeting. Table 2 reports summary statistics for all six series and for their totals.

The median meeting contains four tagged reasons, with the interquartile range running from three to five. Two main facts stand out. Inflationary tags exceed disinflationary tags, and demand motives outnumber supply motives roughly two to one, placing demand at the center of most discussions. Heavy disinflationary supply talk—seen during oil-price crashes or the post-Covid unwind—is relatively rare in comparison.

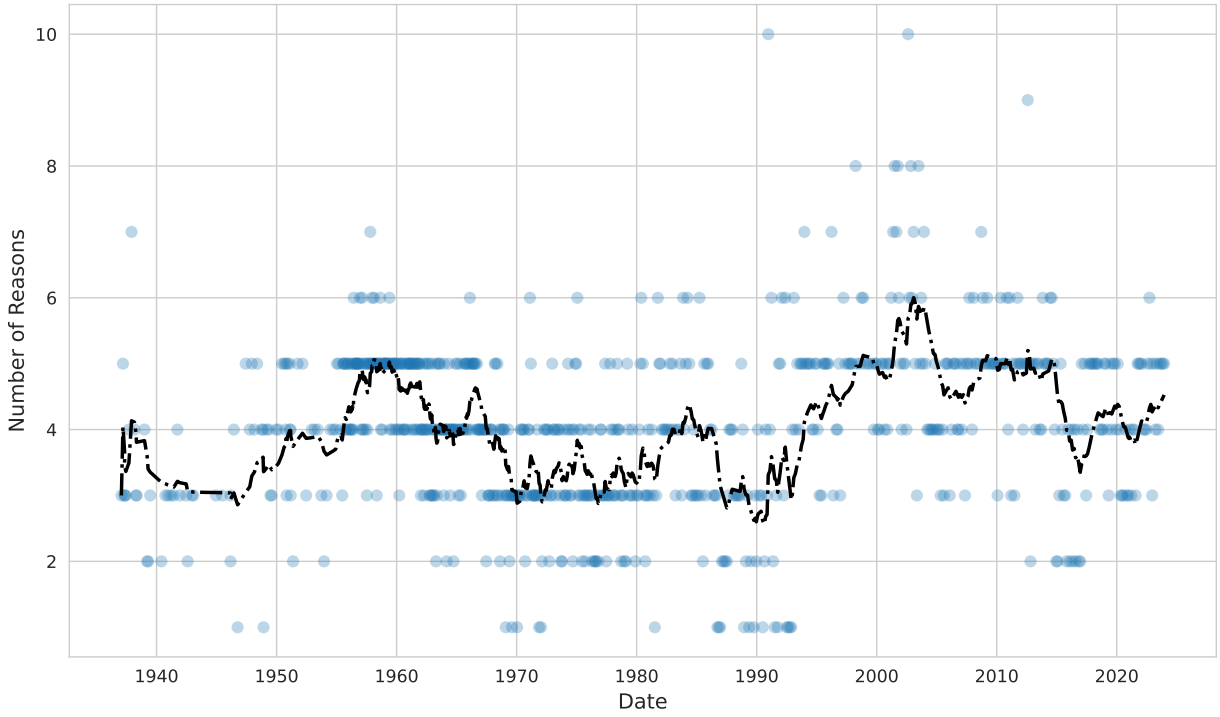
Figure 1 plots, for each meeting, the total number of reasons we extract—pooling demand, supply, and "other," as well as inflationary and disinflationary directions. Dots mark individual meetings; the black curve is an exponential moving average to capture longer trends. Meetings almost always contain three to six distinct reasons and that range is remarkably stable across the entire sample.

Table 2: Summary Statistics for Reason Effects

component	direction	count	mean	std	min	25%	50%	75%	max
Demand	−	776	1.023	1.488	0	0	0	2	8
	+	776	1.434	1.372	0	0	1	2	6
Supply	−	776	0.379	0.706	0	0	0	1	4
	+	776	0.679	0.932	0	0	0	1	5
Other	−	776	0.168	0.446	0	0	0	0	3
	+	776	0.365	0.601	0	0	0	1	3
Total		776	4.081	1.257	1	3	4	5	10

Notes: For each FOMC meeting ($n = 776$) we count how many *demand*, *supply*, and *other* 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.

Figure 1: Total Reason Counts Over Time

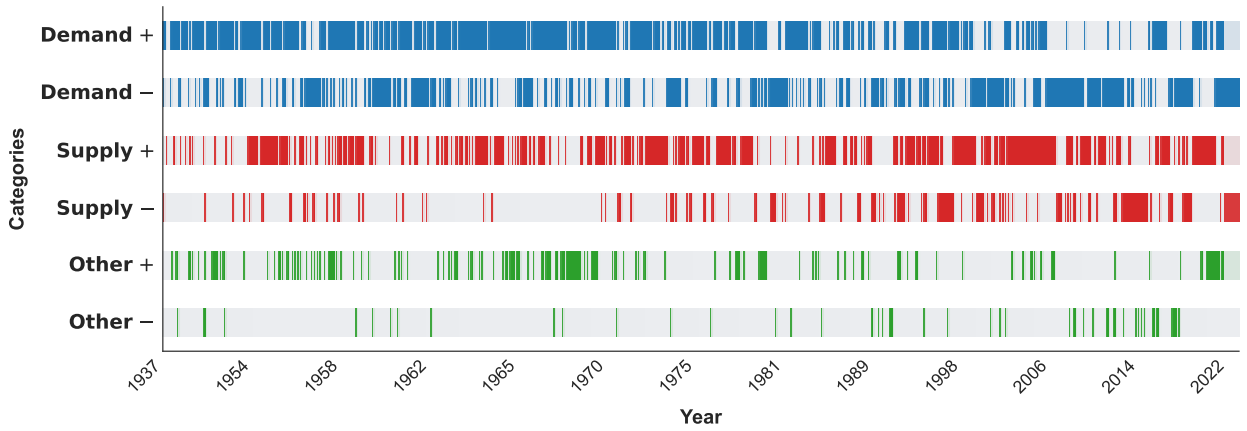


Note: Each dot gives the total meeting-level count of reasons—demand, supply, or other, inflationary or disinflationary. The black curve is an exponential moving average with an eight-meeting half-life, smoothing short-term noise to reveal longer-run trends. See the main text for details on the construction of these variables.

Figure 2 turns to composition. A shaded cell appears whenever a demand-, supply-, or other-type reason is mentioned, giving a time-series "heat map" 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 appear during some of the NBER 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 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).

Figure 2: Presence of Demand and Supply Reasons over Time



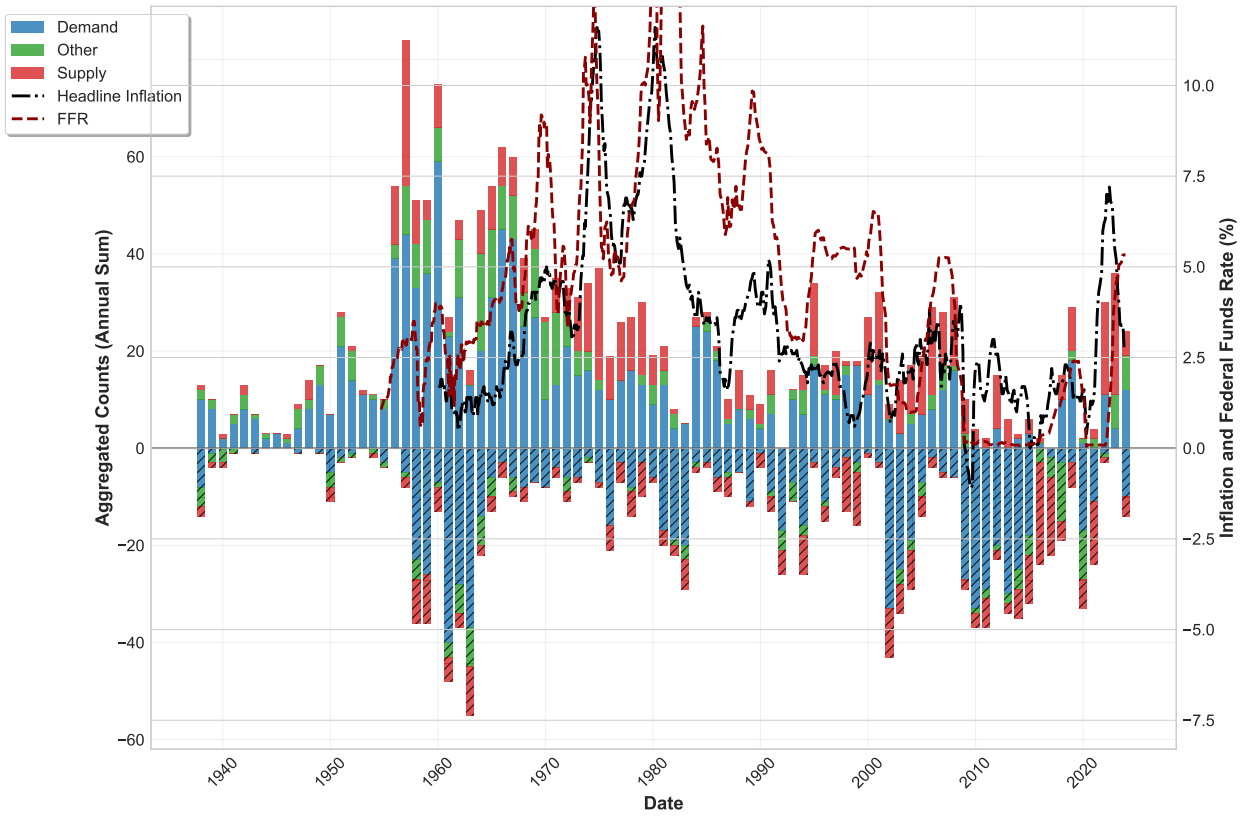
Note: A filled cell marks any FOMC meeting which mentions at least one reason of the indicated category (demand, supply, and other). Inflationary (disinflationary) reasons are accompanied by a "+" ("−"). See the main text for details on the construction of the variables.

2.1 Post-Accord Narratives

Figure 3 shows the annual totals of demand-, supply-, and other-related remarks (left axis). For context, the right axis overlays the federal funds rate and headline inflation. The figure tells a story about some of the key monetary policy episodes in history. Next, we discuss some of these episodes in more detail, relating them to the reasons identified by our methodology.

Eisenhower recession (1957 – 58). During this period, 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

Figure 3: Aggregated Reason Counts Over Time



Notes: This figures aggregates the demand, supply, and "other" 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 (dashed red line) and headline inflation (dash-dotted black line). See the main text for more details on the construction of the demand, supply, and "other" measures from the minutes.

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 brisk 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 demand reasons in 1980, following by a drastic 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 2000s, 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.

Covid (2020 - 2022). The 2020 graph is mostly characterized by disinflationary demand reasons, with a shift in 2021 and 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 also "higher food and energy prices" as well as "Russia's war against Ukraine" as contributing to inflation.

2.2 Net Demand and Supply Contributions

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 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 $x \in \{d, s, o\}$, as

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

The index lies in $[-1, 1]$. 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.

Table 3 reports regressions of inflation on the normalized imbalance measures. Both demand and supply imbalances are positively associated with inflation, but the estimated coefficient on supply imbalance is substantially larger. Moreover, the supply coefficient is statistically significant—at the 1% level for headline inflation and at the 5% level for core inflation—whereas the demand coefficient is small and statistically indistinguishable from zero in both cases.

Table 3: Regression of inflation in imbalance

	Headline	Core
$I_{d,t}$	0.102	0.126
	(0.162)	(0.155)
$I_{s,t}$	0.471***	0.238**
	(0.146)	(0.114)
const	3.203***	3.191***
	(0.238)	(0.217)
Observations	610	610
R^2	0.045	0.019
Adjusted R^2	0.042	0.016

Notes: The dependent variables are headline and core inflation, expressed as year-over-year percentage changes, from the FRED. Each imbalance variable is divided by its sample standard deviation so that the reported coefficients measure the effect of a one-standard-deviation increase in the imbalance. Regressions are estimated by OLS with Newey-West standard errors (in parentheses). *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively.

This pattern aligns closely with standard New Keynesian predictions: monetary policy stabilizes demand-driven fluctuations but cannot fully offset supply shocks. The supply coefficient in the core inflation regression is roughly half that of the headline inflation regression. This is expected, since core inflation excludes food and energy prices.

Overall, our text-based imbalance series closely track monetary history. Supply-side imbalances peak during recognized cost-push episodes, while demand-side imbalances mirror business-cycle fluctuations and reverse sign in downturns as policy counters slack. Both supply and demand imbalances predict subsequent inflation outcomes, with a particularly strong link for supply shocks.

In short, the FOMC minutes and subsequent data tell a consistent story. Having passed this preliminary "smell test", we now test how our decomposition can help explain policy responses over time.

3 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. We then estimate a dynamic Taylor rule that explicitly conditions these responses on the monetary authority's assessment of the underlying drivers of inflation.

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

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

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

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

where g_t and u_t are the shocks to the IS and the PC, respectively. The critical deviation from CGG is that the these shocks have a latent factor structure:

$$g_t = \kappa_d d_t \tag{G}$$

$$u_t = \phi_s s_t, \tag{U}$$

where ϕ_s, $\kappa_d > 0$ are the factor loadings, measuring the differential impact that demand and supply shocks have on the output gap and inflation, respectively. Equivalently, we can

express the factor structure in matrix form:

$$\begin{bmatrix} g_t \\ u_t \end{bmatrix} = \underbrace{\begin{bmatrix} \kappa_d & 0 \\ 0 & \phi_s \end{bmatrix}}_{\Phi} \begin{bmatrix} d_t \\ s_t \end{bmatrix}$$

The baseline CGG framework can be nested under Φ equal to the identity matrix. In this specification we treat d_t and s_t as latent structural factors that reflect the Fed's assessment of the demand- and supply-driven components of inflation.

According to the equations above, a positive demand shock (demand curve shifted to the right) increases the output gap 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, but 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 don't assume the central bank observes these shocks, but rather that it forms estimates about them given the observable data. The central bank knows the DGPs of the true primitive shocks (denoted with stars) to be:

$$d_t^* = (1 - a_d) d_{t-1}^* + a_d \tilde{d}_t \tag{2}$$

$$s_t^* = (1 - a_s) s_{t-1}^* + a_s \tilde{s}_t \tag{3}$$

where $\tilde{d}_t, \tilde{s}_t$ are white noise innovations to the demand and supply shocks, respectively. Since innovations are white noise, a_d, a_s merely act as their mean-preserving spreads. Indeed, an alternative way to express the DGP above is as an autoregressive process where $a_x \times \tilde{x}_t$ are the innovations, for $x = d$, s. To make this equivalence explicit, we denote $\rho_i := 1 - a_i$, for $i = d$, s, and write the DGP as follows:

$$d_t = \rho_d d_{t-1} + \underbrace{a_d \tilde{d}_t}_{\substack{demand \\ innovation}} \tag{d}$$

$$s_t = \rho_s s_{t-1} + \underbrace{a_s \tilde{s}_t}_{\substack{supply \\ innovation}} \tag{s}$$

The central bank does not directly observe any shocks. Instead, it receives signals z_t^s, z_t^d

about their innovations $\tilde{s}_t$, $\tilde{d}_t$ every period:

$$z_t^s = \tilde{s}_t + \eta_t^s \tag{4}$$
$$z_t^d = \tilde{d}_t + \eta_t^d \tag{5}$$

where η_t^i, $i = d$, s are i.i.d. normal.

For simplicity, we assume the central bank perfectly observes d_{t-1}^* and s_{t-1}^* at t. Thus, the only uncertainty is about the contemporaneous innovation, which is resolved in the next period. Under this simplifying assumption, the central bank's best guess about d_t, s_t is just

$$d_t = (1 - a_d) d_{t-1}^* + \alpha_d z_t^d \tag{6}$$
$$s_t = (1 - a_s) s_{t-1}^* + \alpha_s z_t^s \tag{7}$$

The previous equations show that the central bank makes some forecast of d_t^* and s_t^* conditional on its information set and noisy observations. Given d_t and s_t, the best guesses for u_t and g_t can be pinned down. Importantly, we assume that all of these forecasts are common-knowledge (i.e., the market holds the same expectations about d_t, s_t, and everything else). But rather than directly observing all shocks on the economy, this assumption allows for the central bank to be wrong about its diagnosis of the causes of inflation.

As shown in the appendix, the equilibrium interest rate under the baseline discretionary case will be:

$$i_t = \mathbb{E}_t \pi_{t+1} + \Upsilon_s s_t + \Upsilon_d d_t \tag{8}$$

where Υ_s, Υ_d are functions of the deep parameters only. Knowing the decomposition of g and u is important for the central bank because the latent, deep factors d, s are persistent over time. Consequently, figuring out how much of the shock comes from d_t or s_t will affect the central bank's forecast about future inflation and unemployment, and ultimately affect the interest rate path it sets. Importantly, figuring out which part of inflation comes from s_t or d_t will feed into forward expectations in the rational expectations equilibrium via u_t.

The linear-quadratic policy rule implies a stark asymmetry between demand and supply disturbances. Because the real rate exactly offsets the demand shock $g_t = \kappa_d d_t$, 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 shock, $u_t = \phi_s s_t$, is only partially offset. As we show in Appendix B, under standard parameter choices, we have

$$0 < \Upsilon_s < \Upsilon_d, \tag{9}$$

which implies that 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.

4 An Augmented Taylor Rule

The key question, then, is what constitutes the correct model-based measure of the "contribution of demand and supply to inflation." Under the discretionary equilibrium, rational expectations imply that inflation depends solely on supply shocks. Equation (8) provides precisely such a relation, expressing equilibrium inflation as a linear function of s_t, and not of d_t. We exploit this linear structure for identification in what follows.

Our proxy for the central bank's signals on the innovations $\tilde{d}_t$ and $\tilde{s}_t$ are the imbalances $I_{d,t}$ and $I_{s,t}$, defined in Equation (1). Following the DGPs of d_t and s_t, we assume the proportions of inflation attributed to demand and supply, $P_{d,t}$ and $P_{s,t}$, are exponential moving averages of these innovations:

$$P_{d,t} = (1 - b_d) \times P_{d,t-1} + b_d \times I_{d,t} \tag{10}$$

$$P_{s,t} = (1 - b_s) \times P_{s,t-1} + b_s \times I_{s,t} \tag{11}$$

This exponential moving average structure mirrors the DGP of the measurement equations for d_t and s_t from the central bank's perspective, as in Equations (6) and (7).[5]

Finally, we compute our proxies for d_t and s_t as the interaction between these proportions and contemporaneous inflation:

$$d_t = P_{d,t}(b_d) \times \pi_t \tag{12}$$

$$s_t = P_{s,t}(b_s) \times \pi_t \tag{13}$$

In the above expression, we leave both proportions as functions of the exponential smoothing parameters b_d and b_s to make that dependence explicit.

We calibrate $b_d = b_s = 0.2$ throughout most of our analysis. This parameterization implies a persistence parameter of $\rho_d = \rho_s = 0.8$, and a half-life of approximately three FOMC meetings for the process. This implied persistence is similar to the shocks of Clarida, Galí, and Gertler (2000), but the main results do not change significantly when setting different values for these parameters. Figure C.1 presents the time series plots of P_d and P_s, while

[5] In order to be consistent with the model, we also assume that the imbalances we observe reflect both the observed contemporaneous measurement as well as the correction of the vintage measurement to account for the assumption that d^*_{t-1} and s^*_{t-1} are observed at t.

Figure C.2 presents the plots of d and s (in the Appendix).

Our baseline specification for the Taylor rule is given by

$$i_t = i^* + \beta_x x_t + \beta_\pi \pi_t + \varepsilon_t. \tag{14}$$

We adopt this specification for simplicity and clarity rather than realism. It serves as a conservative baseline precisely because it precludes interest rate smoothing. Alternative specifications incorporating interest rate smoothing are explored in our robustness analyses.

We extend the baseline specification in Equation (14) by introducing differential Federal Reserve responses conditional on inflation being primarily demand- or supply-driven. Specifically, we estimate the augmented Taylor rule:

$$i_t = i^* + \beta_x x_t + \beta_\pi \pi_t + \beta_d d_t + \beta_s s_t + \varepsilon_t. \tag{15}$$

We also consider a specification that replaces d_t and s_t with their proportional counterparts, $P_{d,t}$ and $P_{s,t}$, in the preceding equation.

Overall, holding inflation constant, coefficients β_d and β_s capture how the Fed's policy response varies when it perceives demand versus supply factors as predominantly driving inflation. When using $P_{d,t}$ and $P_{s,t}$, these coefficients represent the Fed's incremental responses to the relative contribution of demand or supply to inflation.

Table 4 presents the estimates of this augmented Taylor rule. The first column regresses the federal funds rate on the output gap and inflation, following Equation (14). The second column augments the specification by including the proportions $P_{s,t}$ and $P_{d,t}$. We observe that a higher share of demand-related reasons significantly predicts higher nominal interest rates, while the supply proportion's coefficient is negative, smaller, and statistically insignificant.

The third column substitutes the direct measures d_t and s_t instead of the proportions, as in Equation (14). Once again, the coefficient of s_t is negative, while the coefficient on d_t is positive. This indicates that the Fed responds less to inflation when it believes it is driven by supply, and more when it believes it is driven by demand. Note that the point estimate of the total demand response ($\beta_\pi + \beta_d$) is much larger than one, while the total supply response ($\beta_\pi + \beta_s$) is well below one.

The last two columns regress headline inflation directly on our measures of demand and supply contributions. While $P_{s,t}$ and s_t are significantly associated with higher inflation rates, the effects of $P_{d,t}$ and d_t are smaller and statistically insignificant.

These findings align with the standard macroeconomic intuition formalized in our model. Yet, they are puzzling from the perspective of the standard Taylor rule, which prescribes a uniform response to inflation irrespective of its source. Instead, the results above highlight

Table 4: Dynamic Taylor Rule and Inflation Regressions

	(1) FFR	(2) FFR	(3) FFR	(4) PCEPI	(5) PCEPI
x_t	0.267**	0.087	0.243***		
	(0.104)	(0.116)	(0.086)		
π	1.017***	1.006***	1.322***		
	(0.104)	(0.114)	(0.146)		
$P_{s,t}$		-1.259		2.975***	
		(1.001)		(0.759)	
$P_{d,t}$		2.001***		0.639	
		(0.604)		(0.435)	
s_t			-0.971***		1.375***
			(0.296)		(0.122)
d_t			0.400*		-0.083
			(0.217)		(0.195)
const	1.490***	1.598***	1.133***	2.761***	2.237***
	(0.300)	(0.296)	(0.286)	(0.223)	(0.207)
Obs.	623	623	623	623	623
R^2	0.554	0.598	0.627	0.155	0.380
Adj. R^2	0.553	0.596	0.625	0.152	0.378

Notes: Columns (1)–(3) estimate a dynamic Taylor rule for the federal funds rate (FFR). Column (1) includes the unemployment gap (x_t) and year-over-year inflation (π). Column (2) adds the share of supply-related reasons ($P_{s,t}$) and demand-related reasons ($P_{d,t}$). Column (3) replaces the shares with the decomposed inflation components s_t (supply) and d_t (demand). Columns (4)–(5) switch the dependent variable to headline inflation (PCEPI). Newey–West standard-errors are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

that supply shocks are strongly related to inflation, while demand shocks are not; at the same time, interest rates respond strongly to demand shocks, and weakly to supply shocks. These results emphasize the importance of differentiating between inflation drivers when analyzing monetary policy from a descriptive standpoint.

5 Robustness tests and further analyses

In this section, we present additional specifications and robustness checks to further validate and elaborate on our Taylor rule results from Section 4. First, we re-estimate our augmented Taylor rule using the shadow interest rate. Second, we investigate potential shifts in policy responsiveness across different time periods. Third, we benchmark our measures against other empirical decompositions of inflation commonly employed in the macroeconomic literature. Additionally, we incorporate interest rate smoothing into our Taylor rule specification to capture gradual policy adjustments. Finally, we compare our constructed measures with the Federal Reserve's Greenbook forecasts for inflation and unemployment. These analyses provide a thorough assessment of the robustness and general applicability of our conclusions.

5.1 Using the shadow rate

Table 4 covers episodes in which the federal funds rate is constrained by the zero lower bound. To address this, Table 5 re-estimates the augmented Taylor rule using the Wu and Xia (2016) shadow federal funds rate (i_t^s) instead of the observed rate:

$$i_t^s = s^* + \beta_x x_t + \beta_\pi \pi_t + \beta_d d_t \text{ (or } P_{d,t}) + \beta_s s_t \text{ (or } P_{s,t}) + \varepsilon_t. \tag{16}$$

The shadow rate series spans January 1990–February 2022, which is why these regressions have fewer observations than the main regressions reported in Table 4. The data is obtained from the Atlanta Fed.[6]

Column (1) regresses i_t^s on the output gap, x_t, and the inflation rate, π_t. The coefficient on x_t is positive and significant, whereas the coefficient on π_t is positive but not statistically different from zero. Thus, once the zero lower bound binds, the shadow rate reacts primarily to real-activity gaps rather than to headline inflation.

Column (2) augments the baseline with the inflation-decomposition shares $P_{s,t}$ and $P_{d,t}$. After these proportions are included, the coefficients on x_t and π_t shrink toward zero and lose significance. By contrast, both decomposition variables enter with large, positive, and

[6] Available at: https://www.atlantafed.org/cqer/research/wu-xia-shadow-federal-funds-rate.

Table 5: Taylor Rule Regressions using the Shadow Rate

	Dependent variable: Shadow Rate		
	(1)	(2)	(3)
x_t	0.783***	0.185	0.574***
	(0.174)	(0.175)	(0.188)
π_t	0.378	0.063	0.604
	(0.468)	(0.522)	(0.554)
$P_{s,t}$		1.898**	
		(0.911)	
$P_{d,t}$		3.903***	
		(0.943)	
s_t			-0.243
			(0.523)
d_t			0.995**
			(0.486)
const	2.128**	2.705***	1.837*
	(0.930)	(0.910)	(0.995)
Obs.	248	248	248
R^2	0.368	0.492	0.411
Adj. R^2	0.362	0.483	0.401

Notes: Columns (1)–(3) estimate a dynamic Taylor rule with the shadow rate from Wu and Xia (2016) as the dependent variable. Column (1) regresses the shadow rate on the unemployment gap (x_t) and year-over-year inflation (π_t) as the baseline specification. Column (2) augments the specification with the shares of supply-related reasons ($P_{s,t}$) and demand-related reasons ($P_{d,t}$). Column (3) replaces these shares with the decomposed inflation components s_t (supply-driven inflation) and d_t (demand-driven inflation), while retaining the baseline controls. Newey–West standard errors are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

statistically significant coefficients, and the regression R^2 rises by more than 10 percentage points. Consistent with earlier tables, the estimate on the demand share is roughly twice that on the supply share, confirming that the central bank policy rate responds more to demand-driven than to supply-driven inflation also around the zero lower bound.

Column (3) replaces the proportion measures with the absolute inflation factors d_t and s_t. Controlling for the current inflation rate, the additional response to d_t is positive and significant, whereas the additional response to s_t is negative and indistinguishable from zero. Taken together, the shadow-rate regressions deliver the same message as the main Taylor rule regressions. If anything, the estimates are larger in magnitude over this sample.

5.2 Taylor rule under different subperiods

Monetary-policy history reveals substantial shifts in the Fed's view of the sources of inflation and its understanding of macroeconomics.[7] For instance, Bernanke (2006) highlights a sharp contrast in policy responses before and after the early 1980s. Prior to that decade, energy-price increases fed into wages and long-run inflation expectations, forcing the Fed to impose steep rate hikes that helped trigger the recessions of the 1970s and early 1980s.

Clarida, Galí, and Gertler (2000) likewise 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 6 re-estimates our augmented Taylor rule specification by subsample.

Columns (1)–(3) cover the pre-Volcker era—Chairs Martin (1951–70), Burns (1970–78), and Miller (1978–79).[8] The baseline coefficient on inflation is below unity, consistent with the accommodative stance documented by Clarida, Galí, and Gertler (2000). Adding the proportional measures $P_{d,t}$ and $P_{s,t}$ leaves that estimate essentially unchanged; their own coefficients are positive but small and statistically insignificant. Replacing the proportions with the direct components d_t and s_t yields the same picture: the incremental demand term is positive and marginally significant ($p < 0.10$), the supply term is near zero and insignificant, and the combined response to demand-driven inflation, $\beta_\pi + \beta_d$, remains below one.

Columns (4)–(6) restrict the sample to the Volcker–Greenspan era and yield markedly different results. The baseline coefficient on inflation now exceeds one, indicating an aggres-

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

[8] The regression begins in 1960, when the FRED headline-inflation series becomes available.

Table 6: Taylor Rule Regressions across Monetary-Policy Periods

	Dependent variable: FFR								
	Pre-Volcker (1960 - 1979)			**Volcker–Greenspan (1979 - 2005)**			**Pre-GFC (1960 - 2007)**		
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**
x_t	0.38***	0.37***	0.32***	0.04	-0.21	0.09	0.15	0.15	0.16*
	(0.05)	(0.06)	(0.05)	(0.15)	(0.14)	(0.16)	(0.10)	(0.11)	(0.09)
π_t	0.79***	0.78***	0.71***	1.28***	1.25***	1.48***	0.96***	1.01***	1.20***
	(0.05)	(0.06)	(0.09)	(0.15)	(0.14)	(0.18)	(0.09)	(0.10)	(0.14)
$P_{s,t}$		0.14			-1.65**			-2.51***	
		(1.01)			(0.74)			(0.91)	
$P_{d,t}$		0.10			3.37***			0.78	
		(0.31)			(0.67)			(0.67)	
s_t			0.10			-1.01***			-0.74***
			(0.17)			(0.28)			(0.25)
d_t			0.20*			0.61**			0.24
			(0.11)			(0.25)			(0.24)
const	2.05***	2.01***	2.05***	2.08***	2.27***	2.04***	2.27***	2.53***	1.95***
	(0.14)	(0.20)	(0.14)	(0.54)	(0.42)	(0.47)	(0.31)	(0.28)	(0.29)
Obs.	275	275	275	206	206	206	497	497	497
R^2	0.80	0.80	0.81	0.66	0.75	0.73	0.58	0.61	0.63
Adj. R^2	0.80	0.80	0.81	0.66	0.74	0.73	0.58	0.61	0.62

Notes: Each trio of columns reports three Taylor-rule specifications from Table 4. Columns are grouped by each sub-sample: the pre-Volcker period (1960 - 1979); the Volcker-Greenspan era (1979 - 2005); and the longer pre-GFC subsample considered by Carvalho, Nechio, and Tristão (2021). Newey–West standard errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

sive anti-inflation stance. When $P_{d,t}$ and $P_{s,t}$ are added, both enter with large and statistically significant coefficients: the supply share carries a negative estimate significant at the 5 percent level, whereas the demand share carries a positive estimate significant at the 1 percent level. These estimates are an order of magnitude larger than their counterparts in the pre-Volcker subsample, in line with the regime shift in policy reaction to the sources of inflation at that time.

Column (6) replaces $P_{d,t}$ and $P_{s,t}$ with the absolute contributions d_t and s_t; the results are qualitatively the same. The coefficient of d_t is strongly positive. The coefficient of s_t is strongly negative, and the inflation coefficient remains above unity. These estimates confirm the post-1980 shift from accommodation to activism: the Fed reacts strongly to demand-driven inflation while largely discounting supply-driven inflation. By contrast, the pre-Volcker estimates remain accommodative and exhibit no systematic differentiation between demand and supply sources of inflation.

A comparison of the R^2 between the pre- and post-Volcker periods also reveals a substantial difference. The R^2 of the baseline specification is 80% pre-Volcker and 66% post-Volcker. This finding is consistent with Orphanides (2003), which finds that the Taylor rule's fit as a description of monetary policy is higher in the pre-Volcker period, while the biggest departures to the rule happen during Volcker's disinflation period. Adding our demand and supply

contributions helps closing this gap: while the R^2 in the pre-Volcker period stays roughly unchanged in the augmented rule, the R^2 in the post-Volcker period jumps to around 73% to 75%. The improvement suggests that the Fed's post-1980 policy framework incorporated explicit differentiation between demand- and supply-driven inflation, consistent with a more sophisticated inflation-targeting strategy and a departure of the standard rule.

For completeness, we re-estimate the Taylor rule on the 1960–2007 subsample—spanning the pre-Volcker years through the eve of the Great Financial Crisis—to facilitate comparison with recent studies such as Carvalho, Nechio, and Tristão (2021) and Hofmann, Manea, and Mojon (2024). In the baseline specification, the inflation coefficient is just below unity. After adding $P_{d,t}$ and $P_{s,t}$, the point estimate rises slightly above one in Column (8). The incremental response to $P_{d,t}$ is positive but insignificant, whereas the response to $P_{s,t}$ is large, negative, and highly significant.

Column (9) replaces the proportions with d_t and s_t. The inflation coefficient moves to 1.20; the extra coefficient on d_t is positive but small and insignificant, while that on s_t is -0.74 and significant. Hence, the total response to supply-driven inflation is 0.46, well below unity, whereas the response to demand-driven inflation is 1.44, above unity. The roughly three-to-one responsiveness to demand versus supply is similar to the results in Hofmann, Manea, and Mojon (2024). Unlike that study, however—which allows for interest-rate smoothing—our simpler specification yields a total supply response coefficient below one.

Across regimes, our estimates document a clear shift in Fed behavior. Before 1979, the inflation coefficient was below one and neither $P_{d,t}$ nor $P_{s,t}$ improved the fit, matching the accommodative, undifferentiated stance described by Clarida, Galí, and Gertler (2000). After Volcker's arrival, the coefficient exceeds one, the weight on demand-driven inflation is roughly triple that on supply shocks, and the regression R^2 jumps once the decomposition variables are added. Such patterns are consistent with the departure from the rule discussed by Orphanides (2003) and in line with the "targeted" Taylor-rule evidence reported in recent work by Hofmann, Manea, and Mojon (2024). Together, these results confirm the transition from a passive, one-size-fits-all policy to an active rule that leans hard against demand-side inflation while largely looking through supply shocks.

5.3 Comparison to Shapiro (2024)

Recent work attempts to isolate the demand- and supply-side drivers of inflation from aggregate price data. 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.[9] Shapiro thus produces three time series: dem_t, sup_t and ambig_t. 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 estimatos from Shapiro (2024) comparable to ours, we normalize the absolute demand and supply series and convert them into proportional contributions like ours:

$$\text{dem}_t^* = \frac{\text{dem}_t}{|\text{dem}_t| + |\text{sup}_t| + |\text{ambig}_t|}, \qquad \text{sup}_t^* = \frac{\text{sup}_t}{|\text{dem}_t| + |\text{sup}_t| + |\text{ambig}_t|}. \tag{17}$$

Table 7 presents the regressions of the normalized PCE measures from Shapiro (2024) on our proportional shares.

The top panel uses the variant of the PCE measure that is based on headline inflation. We see that both demand and supply PCE measures load positively on our text–based counterparts, while the cross–effects are small. The bottom panel repeats the exercise with core inflation.

Two differences emerge. First, the R^2 rises for the demand equation but falls for the supply equation. Second, the loading of sup_t^* on $prop_{s,t}$ shrinks sharply. The pattern is intuitive: core inflation removes food and energy, the very items most often cited as supply shocks in the minutes, as seen in Table 1. The evidence therefore suggests that FOMC participants frequently diagnose and discuss headline-level pressures even though core inflation guides its near-term policy, while headline inflation anchors the long-run mandate (Mishkin, 2007).

Table 8 compares our text factors with the PCE factors as additional variables in the Taylor rule. We restrict our series to match Shapiro (2024), which starts in 1969. In these regressions, we directly include the measures supplied by the author without any transformations.

Columns (1) and (2) compare the two measures for headline inflation. In Column (1), we see essentially the same result as reported in Table 4 for our text factors, now in this different subsample. Column (2) shows that the response to dem_t is insignificant when controlling for inflation. In addition, the response has the opposite sign to what we would expect from the theory.

Columns (3) and (4) display the results using core inflation. Column (3) repeats the pattern. The incremental response is smaller against s_t and larger with d_t. The difference in the coefficients of d_t and s_t is still large and significant ($p = 0.014$). Column (4) shows again that the PCE factors are mute; both have negative signs and their difference is insignificant.

For completeness, Figure C.3 in the Appendix plots the fitted values from the baseline

[9] Available at: https://sites.google.com/site/adamshap/data-and-code.

Table 7: Comparing PCE-based and text-based inflation decompositions

Panel A: Headline Inflation

	dem_t^*	dem_t^*	dem_t^*	sup_t^*	sup_t^*	sup_t^*
$P_{d,t}$	0.161**		0.173***	0.071		0.033
	(0.066)		(0.066)	(0.048)		(0.043)
$P_{s,t}$		-0.041	-0.085		0.268***	0.260***
		(0.068)	(0.058)		(0.094)	(0.094)
const	0.313***	0.318***	0.329***	0.453***	0.402***	0.404***
	(0.017)	(0.027)	(0.022)	(0.019)	(0.031)	(0.030)
Obs.	463	463	463	463	463	463
R^2	0.102	0.004	0.117	0.018	0.144	0.148
Adj. R^2	0.100	0.001	0.114	0.016	0.142	0.144

Panel B: Core Inflation

	dem_t^*	dem_t^*	dem_t^*	sup_t^*	sup_t^*	sup_t^*
$P_{d,t}$	0.131***		0.144***	-0.060**		-0.074**
	(0.037)		(0.037)	(0.029)		(0.030)
$P_{s,t}$		-0.053	-0.090**		0.080*	0.099**
		(0.049)	(0.045)		(0.048)	(0.048)
const	0.323***	0.330***	0.339***	0.484***	0.470***	0.466***
	(0.011)	(0.018)	(0.015)	(0.012)	(0.016)	(0.016)
Obs.	463	463	463	463	463	463
R^2	0.147	0.014	0.185	0.037	0.037	0.091
Adj. R^2	0.146	0.012	0.182	0.035	0.035	0.087

Notes: The table presents the regressions of the proportional contributions based on the PCE from Shapiro (2024) defined in Equation (17) on our proportional shares, defined in Equations (10) and (11). The top panel uses the headline inflation and the bottom panel uses core inflation to construct the PCE contributions. Newey-Wests standard-errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1%. levels, respectively.

Table 8: Comparing the PCE-Based and Text-Based Augmented Taylor rules

	Dependent variable: Federal Funds Rate			
	Headline		Core	
	(1)	(2)	(3)	(4)
x_t	0.326***	0.344**	0.522***	0.459***
	(0.125)	(0.147)	(0.119)	(0.140)
π_t	1.386***	0.369	1.498***	2.419**
	(0.148)	(0.786)	(0.160)	(1.143)
s_t	-1.075***		-0.772**	
	(0.296)		(0.299)	
d_t	0.417*		0.200	
	(0.213)		(0.247)	
sup_t		1.265		-1.255
		(0.990)		(1.281)
dem_t		-0.144		-1.494
		(0.998)		(1.645)
Obs.	463	463	463	463
R^2	0.603	0.547	0.627	0.596
Adj. R^2	0.600	0.543	0.624	0.592

Notes: The table presents a comparison between the simple augmented Taylor rule defined in Equation (15) using the PCE-based decomposition from Shapiro (2024) and our text-based measure. Columns (1) and (2) use headline inflation for all measures and Columns (3) and (4) use core inflation for all measures. Newey-Wests standard-errors are in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1%.

Taylor rule in Equation (14) alongside the augmented rules using our text-based measures and the PCE-based measures, all constructed using core inflation. Even in this case, where the PCE measures perform best, the augmented fit is nearly identical to the baseline.

In summary, our text-based demand and supply series align with the PCE decomposition of Shapiro (2024), but also provide additional information. In the regressions of Table 7, each text-derived share loads positively on its PCE counterpart, confirming a common classification of price movements. In the Taylor rule estimates (Table 8) the federal funds rate reacts with the expected signs to the text factors, while the PCE factors are insignificant. The minutes therefore contain information on the sources of inflation that expenditure data alone does not capture.

5.4 Taylor Rule with Interest Rate Smoothing

The specification of the Taylor rule described by Equations (14) and (15) is admittedly simple. However, it is not a realistic description of monetary policy response because it is too restrictive. As pointed out by Clarida, Galí, and Gertler (2000), this specification assumes the Fed reacts immediately to changes in inflation and the output gap.

A more realistic model should incorporate the fact that the Fed tends to make smooth changes in interest rates over time. 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.[10]

We follow Hofmann, Manea, and Mojon (2024) (HMM) and define our baseline generalized specification as:

$$i_t = i^* + \rho i_{t-1} + (1-\rho)\left[\beta_\pi(\pi_t - \pi^*) + \varphi_y y_t\right] + \varepsilon_t. \tag{18}$$

As before, i_t is the policy rate, π_t is inflation, π^* is the inflation target, and y_t is the log output gap.[11]

HMM augment this specification with demand and supply components of inflation using the methodology of Shapiro (2024). We follow their specification and define a Taylor rule

[10] A number of papers have been written on this topic, and an 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.

[11] 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".

that includes our demand and supply components:

$$i_t = i^* + \rho i_{t-1} + (1-\rho)\left[\beta_\pi(\pi_t - \pi^*) + \varphi_d(d_t - d^*) + \varphi_s(s_t - s^*) + \varphi_y y_t\right] + \varepsilon_t. \tag{19}$$

where d^*, s^* are the implied targets of the demand and supply components.

Following the methodology from Carvalho, Nechio, and Tristão (2021), we are able to estimate the structural specification from Equations (18) and (19) by OLS. We can do so by defining the auxiliary reduced-form equation similarly to HMM:

$$i_t = \beta_0^{aux} + \rho i_{t-1} + \varphi_\pi^{aux}\pi_t + \varphi_d^{aux} d_t + \varphi_s^{aux} s_t + \varphi_y^{aux} y_t + \varepsilon_t \tag{20}$$

We first estimate Equation (20) by OLS. Structural coefficients are then obtained as

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

Table 9 reports the results. Each block estimates Equation (19) on a different subsample. Newey–West errors are computed for the auxiliary regression and converted to the structural form with the Delta method.

The first row presents the baseline results of our full sample from 1960 to 2023. The response to inflation is above unity and very similar in magnitude to that of Table 4 using the simple Taylor rule. Although the coefficient on the output gap is positive but insignificant, the incremental response to demand is very large and highly significant. The incremental response to the supply component is large and negative, but statistically insignificant.

For clarity, the final two columns report the total policy response to the demand and supply components, defined as $\varphi_\pi + \varphi_d$ and $\varphi_\pi + \varphi_s$, respectively. The demand-side response is about three—well above the Taylor principle—while the supply-side response is near zero and insignificant.

The next three rows repeat the same exercise for the subsamples considered in Section 5.2: pre-Volcker (1960-1979) and Volcker-Greenspan (1979-2005), as well as the subsample considered by HMM for comparison (1979 - 2007). In the pre-Volcker subsample, the inflation coefficient, φ_π, is below unity and the total response to the demand factor, $\varphi_\pi + \varphi_d$, is 1.4. During the Volcker–Greenspan era φ_π rises above one and the total demand response jumps to 3.25, while the total supply response, $\varphi_\pi + \varphi_s$, remains near zero. These results are roughly in line with Table 6 and the standard current monetary policy view of these two periods discussed in Section 5.2. The HMM window shows a large marginal reaction to demand and a negative, insignificant reaction to supply; the resulting total responses are large for demand (about three) and small for supply (below one and insignificant).

Table 9: Augmented Generalized Taylor Rule Estimates by Different Samples

	φ_π	φ_y	φ_s	φ_d	ρ	$\varphi_\pi + \varphi_d$	$\varphi_\pi + \varphi_s$
Full sample (1960 - 2023)	1.217***	0.532	-1.074	1.762***	0.939***	2.979***	0.143
	(0.361)	(0.328)	(0.810)	(0.529)	(0.019)	(0.705)	(0.605)
Pre-Volcker (1960 - 1979)	0.742***	0.821***	-0.392	0.752	0.934***	1.494***	0.350
	(0.259)	(0.314)	(0.575)	(0.492)	(0.020)	(0.524)	(0.476)
Volcker–Greenspan (1979 - 2005)	1.513***	0.137	-0.762	1.735***	0.865***	3.249***	0.751
	(0.239)	(0.260)	(0.690)	(0.515)	(0.043)	(0.448)	(0.724)
HMM (1979 - 2007)	1.513***	0.129	-0.779	1.741***	0.866***	3.255***	0.734
	(0.238)	(0.246)	(0.645)	(0.512)	(0.042)	(0.443)	(0.685)

Notes: Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (19), for different samples: *Full sample* = 1960Q1–2023Q4; *Pre-Volcker* = 1960Q1–1979Q2; *Volcker–Greenspan* = 1979Q3–2005Q4; *HMM* = baseline period considered by Hofmann, Manea, and Mojon (2024), 1979Q3–2007Q4. φ_π is the response to headline inflation, φ_y to the (log) output gap, and φ_d and φ_s to demand and supply components of inflation; d_t and s_t are defined in Equations (12) and (13), and both constructed with a smoothing parameter of 0.2. ρ is the interest-rate smoothing parameter; the last two columns give the combined response to demand-driven ($\varphi_\pi + \varphi_d$) and supply-driven ($\varphi_\pi + \varphi_s$) inflation. Coefficients are obtained by OLS estimation of the auxiliary reduced-form in Eq. (20) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors from the auxiliary OLS regressions, and then apply the Delta Method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

The coefficients estimated for the baseline sample are similar in magnitude to the post-Volcker coefficients reported by HMM (although their baseline sample is more limited). On the other hand, while HMM estimate a response to the supply component that is larger in the pre-Volcker period, we find a smaller response in that period and an overall small response to the supply component in general (in particular, smaller than unity). However, it must be noted again that the methodology behind the construction of these measures is very distinct. While HMM uses the sign restriction changes in the sub-categories of the PCE from Shapiro (2024) to identify demand and supply, our measure is based on the discussion of the FOMC participants at the meetings. Furthermore, the persistence parameter, $\hat{\rho}$, is large in all samples, in line with the well-documented persistence in the policy rate.

Taken together, the evidence confirms the pattern in the simpler Taylor rule (Table 4). Even after allowing for interest rate smoothing, we find that the Committee reacts strongly to demand shocks, only weakly to supply shocks, and shifts from a passive stance before 1979 to an active, targeted stance after Chair Volcker.

5.5 Greenbook Estimates

Our measure allows us to ask whether the FOMC's inflation forecasts change with the Fed's diagnosis of demand versus supply pressures. The model equilibrium (presented in Appendix B) implies that forecasted inflation should be a linear function of the supply factor only:

$$\mathbb{E}_t[\pi_{t+h}] \;=\; \frac{\alpha\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)}\,\rho_s^{\,h}\,s_t, \qquad h \geq 0.$$

If s_t is stationary, the right-hand side term decays with horizon h and forecasted inflation converges to zero as $h \to \infty$. Under the baseline smoothing $b_d = b_s = 0.2$, however, the factors d_t and s_t are very persistent, with autocorrelations 0.94 and 0.97 (Figure C.2). Expectations should, therefore, inherit this persistence.

We test the model's predictions with Greenbook (more recently, Tealbook) projections of quarterly headline CPI inflation. For each meeting, we take horizons $h = 0$ (nowcast) through $h = 4$. The inflation series begin in 1979 Q4; unemployment forecasts, included for additional estimations, begin in 1967 Q1. Both data sets end in 2019 Q4 owing to the five-year release lag. Table 10 reports the estimates.

The upper panel regresses the inflation forecast for quarter $t+h$ on d_t and s_t. Expected inflation reacts far more to the supply factor than to the demand factor. The s_t coefficient is large at $h = 0$, drops sharply at $h = 1$, and levels off for $h \geq 2$. The d_t coefficient is small, statistically insignificant, and only slightly upward sloping with horizon.

The lower panel regresses unemployment forecasts on d_t and s_t. Demand shocks push

Table 10: Effects of the demand and supply components on the Greenbook forecasts

(a) Inflation forecasts

	Forecast of inflation at quarter $t+h$				
	h=0	h=1	h=2	h=3	h=4
d_t	0.263	0.297	0.305	0.301	0.332
	(0.350)	(0.357)	(0.323)	(0.296)	(0.284)
s_t	1.565***	1.013**	0.503	0.491	0.524*
	(0.495)	(0.513)	(0.360)	(0.312)	(0.311)
const	2.486***	2.732***	2.908***	2.827***	2.799***
	(0.391)	(0.378)	(0.325)	(0.275)	(0.257)
Observations	293	293	293	293	293
R^2	0.257	0.155	0.081	0.102	0.129
Adj. R^2	0.252	0.149	0.075	0.096	0.123

(b) Unemployment forecasts

	Forecast of unemployment at quarter $t+h$				
	h=0	h=1	h=2	h=3	h=4
d_t	−0.454***	−0.461***	−0.433***	−0.404***	−0.379***
	(0.121)	(0.118)	(0.114)	(0.112)	(0.116)
s_t	0.061	0.114	0.156*	0.188**	0.210**
	(0.089)	(0.088)	(0.090)	(0.092)	(0.102)
const	6.154***	6.181***	6.197***	6.186***	6.164***
	(0.221)	(0.222)	(0.222)	(0.223)	(0.232)
Observations	443	434	423	405	375
R^2	0.145	0.148	0.136	0.124	0.111
Adj. R^2	0.141	0.144	0.132	0.120	0.107

Notes: This table reports OLS regressions of Fed's Tealbook (formerly Greenbook) forecasts on the contemporaneous demand- and supply-driven components of inflation measured at each FOMC meeting. The upper panel uses the forecast of headline CPI inflation for quarter $t+h$ as the dependent variable; the lower panel uses the forecast of the unemployment rate for the same horizons. Columns correspond to forecast horizons $h = 0, \ldots, 4$, with horizon $h = 0$ referring to the current quarter ("nowcast"). The explanatory variables are: d_t, the demand-related component of inflation, and s_t, the supply-related component of inflation, both constructed as described Equations (12) and (13). The inflation-forecast sample runs from 1979:Q4 to 2019:Q4 ($N = 293$); the unemployment-forecast sample runs from 1967:Q1 to 2019:Q4, with fewer observations at longer horizons because forward projections are missing in some early vintages. Newey-West standard errors are shown in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1%.

expected unemployment down: the d_t coefficients are large, negative, and roughly flat across horizons. Supply shocks push it up: the s_t coefficients start small, rise with h, and reach significance at the 10 percent level for $h = 2$ and 5 percent for $h = 3$–4. Moreover, and reassuringly, the opposite signs match our identification: demand shocks are expansionary, supply shocks contractionary.

The forecast regressions mirror the contemporaneous estimates in Table 4. The coefficient signs align with our identification: demand shocks increase inflation and lower unemployment, whereas supply shocks raise both. The responses fade only gradually with the horizon, a pattern consistent with the high empirical persistence of d_t and s_t and with the prediction of the model that persistent shocks translate into persistent expectations.

6 Model Calibration

In this section, we estimate the factor loadings we introduced in the model: ϕ_s and κ_d. We estimate the parameters by OLS after making variable transformations. First, note that:

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

We substitute that expression into Equation (37) to obtain a linear equation on d_t and s_t:

$$i_t = \phi_s \left[\frac{\alpha \rho_s + \frac{\lambda}{\psi}(1 - \rho_s)}{\lambda^2 + \alpha(1 - \beta \rho_s)} \right] s_t + \kappa_d \frac{1}{\psi} d_t \tag{22}$$

Next, we define the transformed variables:

$$\tilde{s}_t := \left[\frac{\alpha \rho_s + \frac{\lambda}{\psi}(1 - \rho_s)}{\lambda^2 + \alpha(1 - \beta \rho_s)} \right] s_t \tag{23}$$

$$\tilde{d}_t := \frac{d_t}{\psi} \tag{24}$$

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

We fix the remaining variables at standard values from the literature: $\beta = 0.99$, $\lambda = 0.3$, $\alpha = 0.048$, and $\psi = 1$.[12] We set $b_d = b_s = 0.2$ as before. This implies $\hat{\rho}_s = 0.97$ and $\hat{\rho}_d = 0.94$.

[12] The values for ψ, λ, and β are from Clarida, Galí, and Gertler (2000). The value for α comes from Woodford (2009) and is also considered by Debortoli et al. (2018) in their analysis of the weights in the loss function of the central bank.

We also use Newey-West standard errors.[13] Table 11 presents the results.

Table 11: Model Estimation

	Estimate	s.e.	t-stat
ϕ_s	0.0758	0.020	3.740
κ_d	0.6336	0.312	2.030

Notes: The table reports OLS estimates of the factor loadings κ_d and ϕ_s from Equation (22). We use the Wu and Xia (2016) shadow rate whenever it is available and the federal funds rate otherwise. s_t and d_t are constructed with headline inflation. The calibrated parameters are: $\beta = 0.99$, $\lambda = 0.03$, $\alpha = 0.048$, $\psi = 1$, and $b_d = b_s = 0.2$. Standard errors are Newey–West. See the main text for more details about the estimation and the choice of the calibrated parameters.

From the table, we see that both parameters have the expected positive sign and are statistically significant. This result is in line with the reduced-form estimates from Tables 3 and 4 (for contemporaneous inflation) and Table 10 (for forecasted inflation).

Using the estimated coefficients, we compute the model-implied impulse responses. Figure 4 displays three sets of IRFs: the equilibrium policy rate i_{t+h} (top panel), expected inflation $\mathbb{E}_t \pi_{t+h}$ (middle), and expected output gap $\mathbb{E}_t x_{t+h}$ (bottom). The red line shows a unit demand shock (d_t); the blue line shows a unit supply shock (s_t).

Two main facts stand out. Interest rates react more strongly to a demand shock than to a supply shock, yet they also converge back to baseline relatively more slowly after a supply shock. By construction, expected inflation and output gap do not respond to demand shocks in equilibrium; we can see, however, that supply shocks generate deviations that reverse very slowly over time.

These impulse-response functions are consistent with the results presented so far. Interest rates respond more to demand shocks than to supply shocks, while (and partially because) supply shocks have a greater and more persistent effect on inflation than demand shocks.

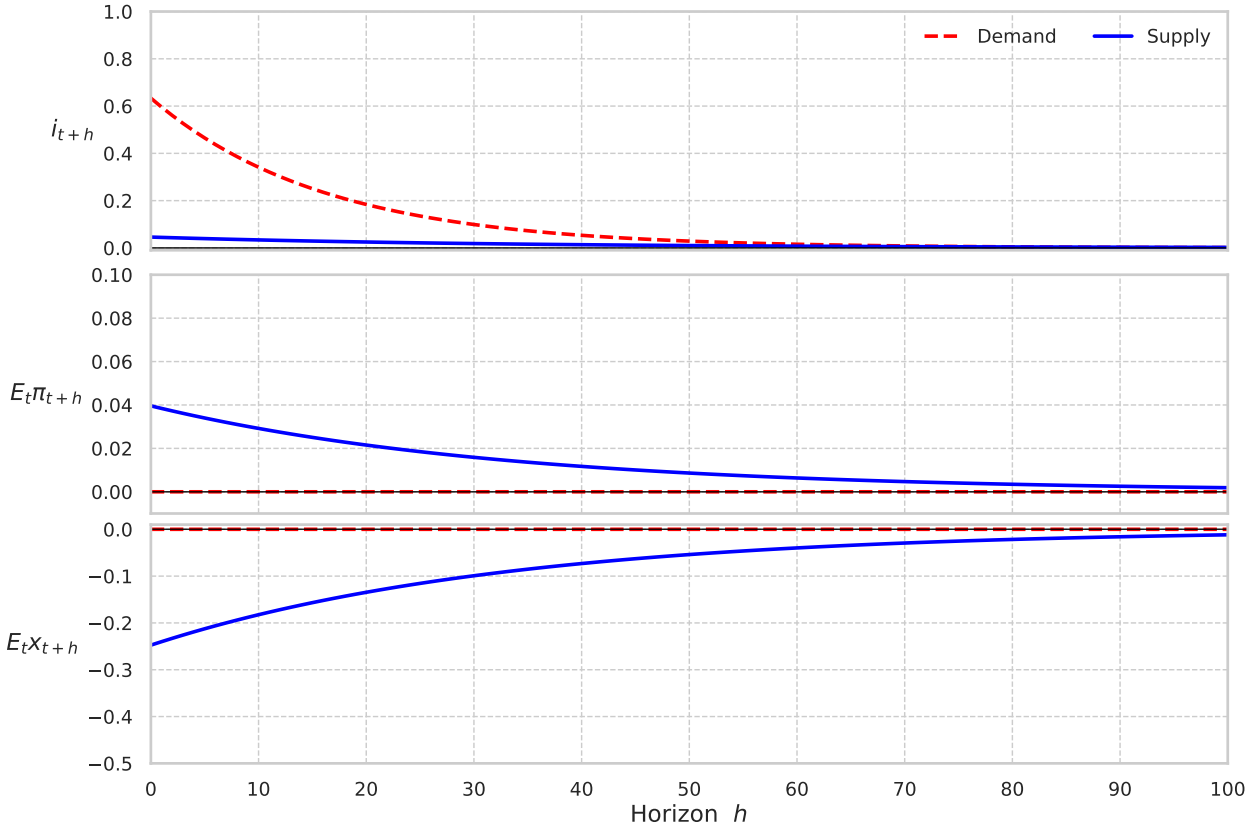
7 Monetary Policy Shocks

Monetary economists track surprises by watching how futures and bond yields jump in the half-hour around an FOMC statement. We take the event-study dataset of Bauer and Swanson (2023*b*) and ask a simple question: when the Fed discusses inflation as demand-driven or supply-driven, do those labels help explain the market's moves?

We measure each monetary policy surprise as the asset's 30-minute return—10 minutes before, 20 minutes after the FOMC announcement—across nine key markets. FF1 and

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

Figure 4: Impulse-Response Functions from the Calibrated Model



Notes: The figure plots impulse response functions implied by the model from Section 3. Using the estimated shocks' loadings ($\hat{\phi}_s$, $\hat{\kappa}_d$), calibrated parameters $\beta, \lambda, \alpha, \psi, b_d, b_s$, and the implied persistences ρ_d, ρ_s, we simulate the equilibrium paths of the policy rate i_{t+h} (top panel), expected inflation $\mathbb{E}_t[\pi_{t+h}]$ (middle), and expected output gap $\mathbb{E}_t[x_{t+h}]$ (bottom) following a one-unit innovation in the demand factor d_t (red dashed line) or the supply factor s_t (blue solid line). Estimation details and parameter values are provided in the main text.

FF2 are the current- and next-month federal-funds futures; ED1 through ED4 are the first four quarterly Eurodollar contracts. We also include longer rates: two-, five-, ten-, and thirty-year yields from note- and bond-futures quotes (T2, T5, T10, and T30). Finally, we study the effects on equity markets: we used scaled log changes in the two S&P 500 futures contracts—the original contract and the later e-mini—by 100, so a one-point move equals one basis point. To capture volatility, we add the log change in the VIXY short-term VIX futures ETF.[14]

As discussed in the introduction, many influential papers in the monetary policy literature have used some subset of the series of asset prices that we consider above. Kuttner (2001) use a scaled version of the FF1 as his monetary policy shock. Nakamura and Steinsson (2018) and Gürkaynak, Sack, and Swanson (2005) use scaled versions of FF1 and FF2, along with ED2, ED3, and ED4. Swanson (2021) include S&P 500 futures, FF1, FF2, ED1 to ED4, T2, T5, T10, and T30, among other series.

Bauer and Swanson (2023*b*) construct their monetary policy surprise (MPS) using ED1 to ED4 around an extended sample of announcements. They additionally extend the analysis from Bauer and Swanson (2023*a*) and construct an orthogonal surprise ($\text{MPS}^{\perp}$) by regressing MPS on a series of macroeconomic and financial predictors prior to the announcement. We consider these two series separately as well.

Table 12 presents the regression results of the high-frequency changes described above on our standardized imbalance measures. We use the raw imbalance series, not the smoothed proportional or level factors. These imbalances are empirical analogue of the innovations in Equations (10) and (11), making them the right choice for our high-frequency regressions.

MPS responds positively and significantly to demand and supply imbalances. When using $\text{MPS}^{\perp}$, point-estimates are positive, but insignificant. Equity market responses are intuitive. Our imbalances are associated with higher volatility and lower market returns, yet only the supply term is statistically significant at 1% for VIXY.

Both measures are positively related to Eurodollar futures and Treasury yields, but their effects vary across maturities. The supply coefficient is larger for the one quarter ahead Eurodollar contract, but becomes smaller and only marginally significant at longer maturities. The demand coefficient is more important at the longer end of the yield curve, showing large, significant effects on three- and four-quarter Eurodollars and on two- and five-year Treasuries.

So far, most of the literature on monetary policy surprises has used the market's response

[14]All data used available at the San Francisco Fed website, except for the VIXY data, which we obtain directly from the Trades and Quotes (TAQ) database. As in Bauer and Swanson (2023*b*), we use the "big" S&P 500 futures up to 1997, after which we use the smaller, more liquid e-mini contract. See also Swanson and Jayawickrema (2023) for more details about the construction of the asset price changes.

Table 12: MPS Regression Results

	MPS	MPS$^{\perp}$	VIXY	S&P500	S&P500 e-mini	FF1	FF2
$I_{d,t}$	0.010**	0.000	0.002	-0.039	-0.010	-0.001	-0.003*
	(0.004)	(0.004)	(0.002)	(0.038)	(0.040)	(0.001)	(0.002)
$I_{s,t}$	0.006**	0.002	0.003***	-0.034	-0.030	0.001*	0.002**
	(0.003)	(0.002)	(0.001)	(0.025)	(0.026)	(0.001)	(0.001)
Obs.	230	223	103	198	205	112	112
R^2	0.047	0.003	0.070	0.012	0.006	0.026	0.039
Adj. R^2	0.039	-0.006	0.052	0.002	-0.004	0.008	0.022
F	5.536***	0.378	4.411**	1.507	0.764	3.460**	5.667***

	ED1	ED2	ED3	ED4	T2	T5	T10	T30
$I_{d,t}$	0.002	0.009**	0.014***	0.015***	0.011***	0.010**	0.008**	0.006*
	(0.003)	(0.004)	(0.005)	(0.005)	(0.004)	(0.004)	(0.004)	(0.003)
$I_{s,t}$	0.005***	0.006**	0.005*	0.005	0.005*	0.005*	0.004*	0.003
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
Obs.	230	230	230	230	230	230	230	230
R^2	0.025	0.048	0.059	0.051	0.055	0.047	0.047	0.028
Adj. R^2	0.016	0.040	0.050	0.043	0.047	0.038	0.038	0.019
F	4.035**	5.562***	6.659***	6.017***	6.803***	5.263***	4.967***	2.943*

Notes: Regression results of high-frequency asset price changes around FOMC announcements on our standardized imbalance measures defined in Equation (1). The time interval is between 10 minutes before the announcement and 20 minutes after the announcement. FF1 and FF2 are the current- and next-month federal-funds futures; ED1 through ED4 are the first four quarterly Eurodollar contracts; T2, T5, T10, and T30 are the two-, five-, ten-, and thirty-year yields from note- and bond-futures quotes; S&P 500 and S&P 500 emini are 100 times the log changes in the big and small version of the contract; VIXY is the log change of the VIX futures ETF. All data on the dependent variables comes from Bauer and Swanson (2023*b*), except for VIXY, which comes from the Trades and Quotes (TAQ) database. See Bauer and Swanson (2023*b*) and Swanson and Jayawickrema (2023) for more details on the construction of the variables. Newey-West standard erorrs in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

to the shock as the measure of the surprise. New work, ours included, tries to understand the source of this surprise. Taken together, the results above suggest that the Fed's own assessment of inflation can help us understand their policy responses and the pattern of market reaction to them.

8 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 specification with our separately identified demand and supply factors reveals that the federal funds rate reacts far more aggressively to demand-related inflation than to supply-related inflation. The asymmetry is muted before the Volcker disinflation, grows markedly during the Volcker–Greenspan era, and persists thereafter — mirroring previous narrative accounts of the Fed's post-1980 commitment to pre-empt demand overheating while "looking through'' adverse supply episodes.

Regressing the classic high-frequency monetary-policy surprise factors on contemporaneous demand and supply attributions shows that part of the asset-price jumps around FOMC announcements can be explained by the Committee's own diagnosis of inflationary pressures. Incorporating textual information therefore improves our understanding of what the Fed intended when it delivered an action that surprised the market.

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

Bauer, Michael D. and Eric T. Swanson. 2023*a*. "An Alternative Explanation for the "Fed Information Effect"." *American Economic Review*, 113(3): 664–700.

Bauer, Michael D. and Eric T. Swanson. 2023*b*. "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." Proceedings, *Jackson Hole Economic 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.* [Erscheinungsort nicht ermittelbar]:W. W. Norton Company, Incorporated. Description based on publisher supplied metadata and other sources.

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

Bernanke, Ben S. 2006. "Energy and the Economy." *Remarks before the Economic Club of Chicago*, Council of Economic Advisers co-sponsorship.

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:Princeton University Press.

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

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.

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.

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

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.

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*.

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

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*.

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

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

Mishkin, Frederic S. 2007. "Headline versus core inflation in the conduct of monetary policy : a speech at the Business Cycles, International Transmission and Macroeconomic Policies Conference, HEC Montreal, Montreal, Canada, Oc." Board of Governors of the Federal Reserve System (U.S.) Speech 332.

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

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.

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

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.

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. 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.

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

Swanson, Eric T. and Vishuddhi Jayawickrema. 2023. "Speeches by the Fed Chair Are More Important than FOMC Announcements: An Improved High-Frequency Measure of U.S. Monetary Policy Shocks." Working Paper, University of California, Irvine.

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. Includes bibliographical references (pages 747-764) and index.

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

A 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 4o-mini API for this task.[15] In addition, 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 the `effect` can either be `I` or `D`, that `reason` needs to be a string, and so on.

First Stage Prompt: Identifying Reasons

```
You are an economist analyzing the following FOMC minute. 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 effect, use "I" for Increasing or "D" for Decreasing (or "na" if
unclear). The effect should capture whether the factor is perceived by FOMC
participants/Fed as contributing to an increase or a decrease in inflation, all
else constant. For example, if the inflation gap is positive but there is a
negative demand shock perceived by the Fed, that shock should be classified as
contributing to a decrease in inflation.

Answer the following:

Reasons for the Deviation
- List ALL reasons that FOMC participants explicitly cite in the text as
structural causes or underlying factors driving inflation (i.e., only those the
Fed itself points to as causes of inflation). Do not include self-referential
answers such as "increasing inflation" or "inflation concerns".
- For each reason, include:
  -> "reason": A brief name (e.g., "supply chain issues"). If there are no
  reasons, answer "na".
```

[15] More information about the model and the API are available at https://platform.openai.com/docs/models/gpt-4o-mini.

```
  -> "importance_score": A score from 0 to 10. Higher scores are more critical
  and contribute to inflation more than lower scores.
  -> "effect": "I" for Increasing or "D" for Decreasing (or "na" if unclear),
  capturing whether the factor is seen as contributing to an increase or
  decrease in inflation, all else constant.
  -> "explanation": A brief explanation referencing the text that must include a
  direct quote from one of the passages to support this explanation. When
  including direct quotes, wrap them in single quotes (') or escape any double
  quotes so that the JSON remains valid.
- 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>,  "effect": "<I/D/na>", "explanation": "<brief explanation referencing the
text with a direct quote>" } ], "all_reasons": "<Y/N>" }

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

Example Output:

{ "reasons": [ { "reason": "energy shock", "importance_score": 5, "effect": "I",
"explanation": "Minutes noted 'energy supply issues' led to price increases." },
{ "reason": "excess liquidity", "importance_score": 3, "effect": "I",
"explanation": "Abundant liquidity pushed demand higher, as stated 'excess
liquidity' was observed." } ], "all_reasons": "Y" }

Now, here is the FOMC minute to analyze:
```

A.2 Second Stage Prompt

The second step 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 use the GPT o4-mini API. In addition, 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 for Demand, Supply, or Other, based on the sign restriction implied by the inflation and output effects.
4. `scratch:` A causal chain linking the shock with inflation and output (if such a link exists). This serves a dual purpose: it helps the LLM reason consistently and it allows the econometrician to verify the validity of the reasoning, allowing transparency. Indeed, a literature in machine learning has shown that intermediary reasoning significantly help 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 cases, as well as contrasting examples (i.e., wage increased caused by demand versus supply).

Second Stage Prompt: Categorizing Reasons

```
You are an economist tasked with classifying the reason cited by the Federal
Reserve for inflation.
Your goal is to categorize the reason into one of three macroeconomic types:
demand-driven, supply-driven, or other/unclear.

The classification should follow the sign-restriction criterion derived from
standard macroeconomic models and empirical identification strategies:

- A demand shock moves inflation and real activity (e.g., GDP, output,
employment) in the same direction.
- A supply shock moves inflation and real activity in opposite directions.
```

Use this logic to reason about whether the stated cause is consistent with an increase in inflation that is likely accompanied by:
- rising output/employment → demand
- falling output/employment → supply

Use only the information below and this logic, and not information that you have access to about what happened in the future for this exercise.
For each reason, assign it to exactly one of the following three categories:

———————————————————— CATEGORIES
————————————————————

Categories:

1. Supply
   - Supply-side or cost-push shocks.
   - Indicators: Increases in production costs, energy prices, input bottlenecks, productivity changes, natural disasters, or regulation that constrains production.
   - Examples: “Higher oil prices,” “COVID-related supply chain disruptions,” “Agricultural output drop,” “Reduced labor force participation,” “Semiconductor shortage”, “Strike reducing output.”

2. Demand
   - Shocks that influence the aggregate demand for goods and services.
   - Indicators: Surges in household or business spending, fiscal transfers, sentiment changes, credit booms, looser financial conditions stimulating spending, increase in wages due to booming economy.
   - Examples: “Consumer spending rebound,” “Strong labor income,” “Booming housing demand,” “Pent-up demand post-COVID,” “Low interest rates spurring investment.”

3. Other
   - Use this only if:
     (a) The reason is tautological or restates inflation itself (e.g., “inflation is high because prices are rising”), without specifying causes.
     (b) The statement is vague, non-causal, or uses ambiguous macroeconomic language.

```
     (c) The sign of the shock's effect on inflation and output is unclear or
     contradictory.

Important: If a reason clearly matches a demand or supply mechanism, never
assign it to "Other." For example, "rising energy prices" is always supply.

———————————————————— EDGE-CASE OVERRIDES
————————————————————

Some additional edge cases:
• Labor & wages
      Higher wages, greater hiring, or tight labor markets ⇒ **Demand**,
      unless explicitly linked to strikes, lockdowns, or other constraints that
      restrict production
      (then ⇒ **Supply**).

• Trade balance / BoP deficits
      "Trade deficit," "balance-of-payments deficit," or gold inflows ⇒
      **Demand**
      (domestic absorption exceeds output),
      unless the text ties them to higher import costs or shortages that curb
      production
      (then ⇒ **Supply**).

———————————————————— CONTRASTING EXAMPLES
————————————————————

• "Higher nominal wage growth boosting household spending"  → Demand
• "Work stoppages due to strikes limiting output"            → Supply
• "COVID lockdowns reducing labor availability"              → Supply
• "Rising trade deficit financed by capital inflows"         → Demand

———————————————————————————————————— OUTPUT
————————————————————————————————————

Return your classification in strict JSON format, without any additional
explanation, markdown, or comments:

{
```

```
  "scratch": "<brief causal chain, max 40 chars, e.g. 'BoP ⇒ ↑spending ⇒
  ↑π,↑y'>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Other>"
}
- scratch: your step-by-step reasoning (keep it short).
- output_effect: expected direction of real activity.
- inflation_effect: expected direction of prices.
- category: final label.

Below, you will find the reason to classify and the surrounding context from the
FOMC minutes.
```

B Model Equilibrium

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 of d_t^*, s_t^* (d_t, 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 ($\star$) given the Phillips curve (PC).
2. Then, given x_t and π_t, it pins down i_t from the IS 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:

$$\mathscr{L} = -\frac{1}{2}\left[\alpha x_t^2 + \pi_t^2\right] + \gamma(\pi_t - \lambda x_t - f_t) + F_t \tag{25}$$

where

$$f_t = \beta \mathbb{E}_t \pi_{t+1} + u_t \tag{26}$$

$$F_t = -\frac{1}{2}\mathbb{E}_t \left[\sum_{j=1}^{\infty} \beta^j (\alpha x_{t+j}^2 + \pi_{t+j}^2) \right] \tag{27}$$

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

$$x_t = -\frac{\lambda}{\alpha}\pi_t \tag{28}$$

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

$$\pi_t = -\frac{\lambda^2}{\alpha}\pi_t + \beta\mathbb{E}_t\pi_{t+1} + u_t \tag{29}$$

$$\implies \pi_t = \frac{\alpha}{\alpha+\lambda^2}\beta\mathbb{E}_t\pi_{t+1} + \frac{\alpha}{\alpha+\lambda^2}u_t \tag{30}$$

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:

$$\pi_t = \frac{\alpha}{\alpha+\lambda^2}\beta\mathbb{E}_t\pi_{t+1} + \frac{\alpha}{\alpha+\lambda^2}u_t \tag{31}$$

$$\implies \pi_t = \frac{\alpha}{\alpha+\lambda^2}\beta\mathbb{E}_t\pi_{t+1} + \frac{\alpha}{\alpha+\lambda^2}(\phi_s s_t) \tag{32}$$

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

$$\begin{aligned}
\pi_t =& \frac{\alpha}{\alpha+\lambda^2}(\phi_s s_t) + \left(\frac{\alpha}{\alpha+\lambda^2}\right)^2 \beta(\phi_s\rho_s s_t) \\
&+ \left[\frac{\alpha}{\alpha+\lambda^2}\right]^3 \beta^2(\phi_s\rho_s^2 s_t) + \ldots \\
=& \frac{\alpha}{\alpha+\lambda^2}\phi_s\left[s_t + \frac{\beta\alpha\rho_s}{\alpha+\lambda^2}s_t + \left(\frac{\beta\alpha\rho_s}{\alpha+\lambda^2}\right)^2 s_t + \ldots\right] \\
=& \frac{\alpha\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)}s_t
\end{aligned} \tag{33}$$

The expression above decomposes inflation in terms of latent demand and supply factors. From this expression, it also follows that:

$$\mathbb{E}_t\pi_{t+h} = \frac{\alpha\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)}\rho_s^h s_t, \quad \forall h \geq 0 \tag{34}$$

$$x_t = -\frac{\lambda\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)}s_t \tag{35}$$

$$\mathbb{E}_t x_{t+h} = -\frac{\lambda\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)}\rho_s^h s_t, \quad \forall h \geq 0 \tag{36}$$

where the equations above follow from the FOC and the DGPs from Equations (d) and (s).

Notice that the total IS shock (g_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:

$$\begin{aligned} i_t =& \mathbb{E}_t\pi_{t+1} + \frac{1}{\psi}\left[\mathbb{E}_t x_{t+1} - x_t\right] + \frac{1}{\psi} g_t \\ \implies i_t =& \mathbb{E}_t\pi_{t+1} + \frac{1}{\psi}\left[\frac{\lambda\phi_s(1-\rho_s)}{\lambda^2+\alpha(1-\beta\rho_s)}\right] + \frac{1}{\psi}\kappa_d d_t \\ =& \mathbb{E}_t\pi_{t+1} + \underbrace{\frac{1}{\psi}\left[\frac{\lambda(1-\rho_s)\phi_s}{\lambda^2+\alpha(1-\beta\rho_s)}\right]}_{\Upsilon_s} s_t + \underbrace{\frac{1}{\psi}\kappa_d}_{\Upsilon_d} d_t \end{aligned} \tag{37}$$

The equation above implies that the real interest rate ($i_t - \mathbb{E}_t\pi_{t+1}$) will fully offset the demand shock $g_t = \kappa_d d_t$. This can be seen by the last term in i_t, which depends on g_t/ψ. 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 g_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 output 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 keeping the output gap fixed.

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:

$$\lambda^2 + \alpha(1-\beta\rho) - \lambda(1-\rho) > 0 \tag{38}$$

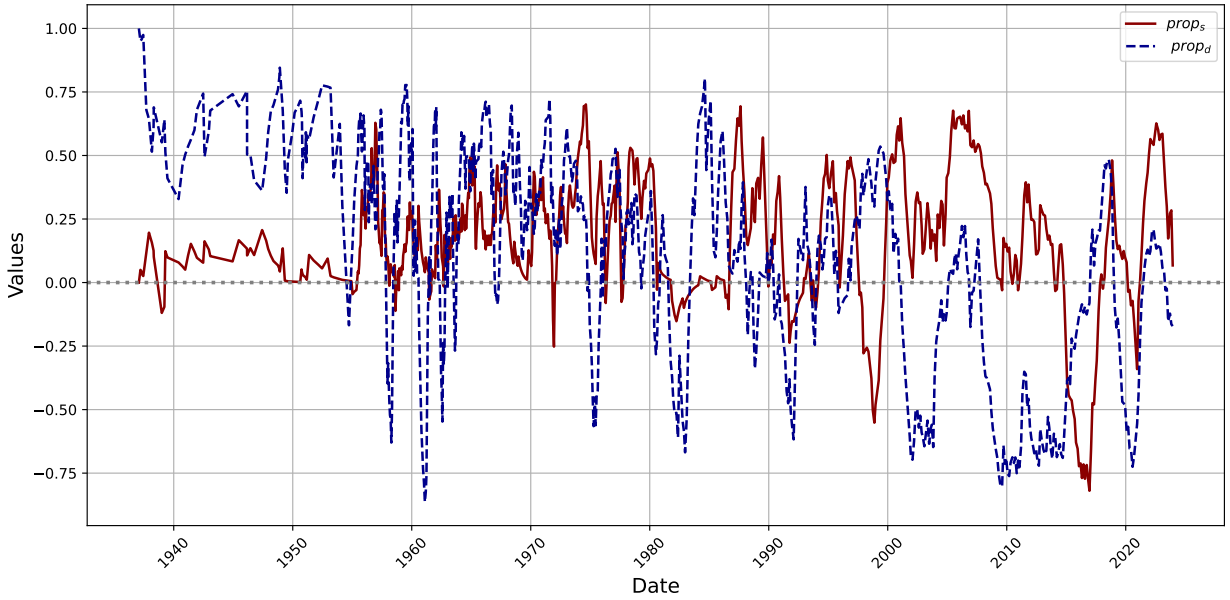
A sufficient condition for the above inequality is $\alpha > \lambda$. This is true with $\lambda = 0.3$ and $\alpha = 0.048$, which are the parameters used by Clarida, Galí, and Gertler (2000) and Woodford (2009), respectively. Thus, in general:

$$0 < \Upsilon_s < \Upsilon_d \tag{39}$$

In summary, everything 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.

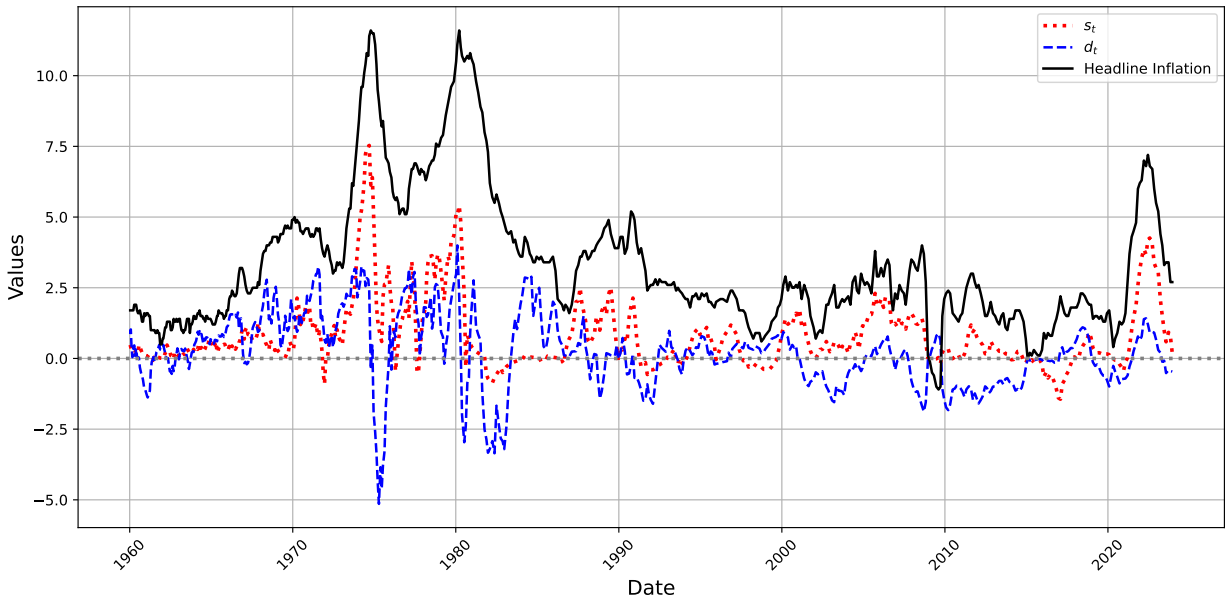
C Supplementary Figures

Figure C.1: Time series of $P_{d,t}$ and $P_{s,t}$



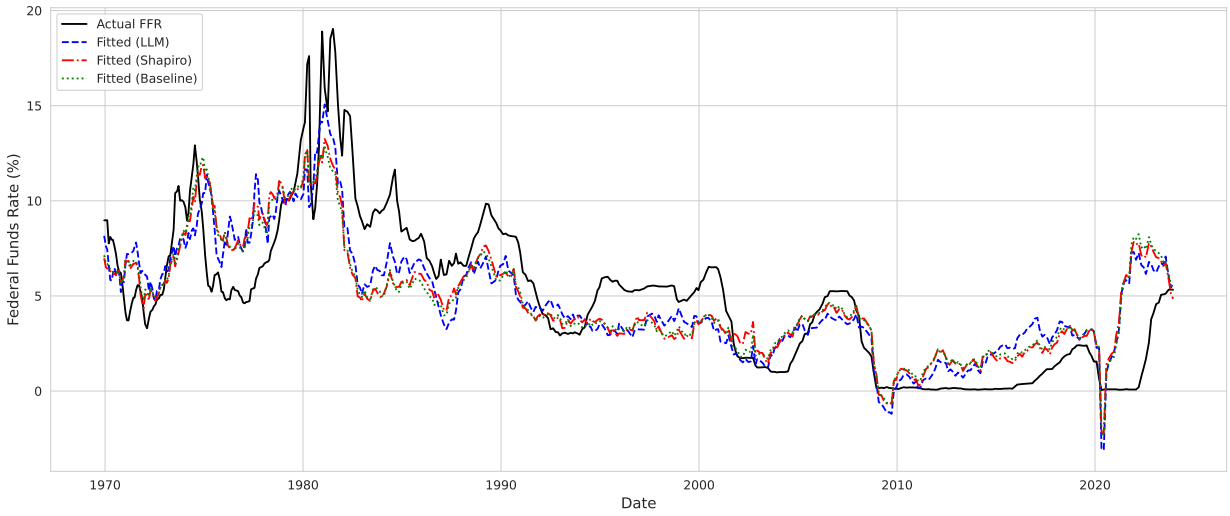
Notes: This plot shows the time-series for our measures of the proportional importance of the demand (blue solid line) and of the supply (red dashed line) components of inflation, $P_{d,t}$ and $P_{s,t}$. These measures are constructed from the imbalances of the counts of demand and supply reasons from the minutes, and defined in Equations (10) and (11). We use a calibrated smoothing parameter of 0.2 in both cases.

Figure C.2: Time series of d_t and s_t



Notes: This plot shows the time-series for our measures of the absolute contributions of the demand (blue dashed line) and of the supply (red dotted line) components of inflation, d_t and s_t. These measures are constructed proportional contributions and defined in Equations (12) and (13). We use a calibrated smoothing parameter of 0.2 in both cases. Additionally, the black solid line displays year-on-year headline inflation (PCE) for the period.

Figure C.3: Comparison of Taylor Rule Fits (Core Inflation)



Notes: The figure plots, for each FOMC meeting from 1969,Q1 to 2023,Q4, the actual federal funds rate together with three sets of fitted values. The first series corresponds to the baseline Taylor rule in Equation (14), using core PCE inflation and the log output gap. The second series adds the text–based demand and supply factors d_t and s_t; the third adds the PCE–based factors dem_t^* and sup_t^* based on Shapiro (2024) following the normalization in Equation (17). See the main text for more details on the construction of the variables.

D 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 Semptember 16th, 2008 meeting. The second comes form the June 15th, 2022 meeting.

Output Example 1: September 2008


```
  {
  "Date": "2008-09-16 00:00:00",
  "all_reasons": "Y",
  "reason": "declining consumer spending",
  "importance_score": 6.0,
  "inflation_effect": "Down",
  "explanation": "The minutes mention that 'consumer spending softened
  appreciably in recent months', contributing to economic weakness.",
  "scratch": "↓cons spend ⇒ ↓y, ↓π",
  "output_effect": "Down",
  "category": "Demand"
}
```


Output Example 2: June 2022


```
 {
  "Date": "2022-06-15 00:00:00",
  "all_reasons": "Y",
  "reason": "COVID-related lockdowns",
  "importance_score": 6.0,
  "inflation_effect": "Up",
  "explanation": "Participants indicated that 'COVID-related lockdowns in
  China were likely to exacerbate supply chain disruptions.'",
  "scratch": "lockdowns ⇒ supply disruption ⇒ ↑π,↓y",
  "output_effect": "Down",
  "category": "Supply"
}
```