Fear of full employment: Labor and inflation at the Fed

Monica DiLeo
Hertie School

Benjamin Braun
London School of Economics

Jérôme Deyris
Sciences Po

This Version: February 2026

Abstract Does the delegation of monetary policy to independent central banks help grow the cake for all, or does it institutionalize a monetary order that allocates a larger slice to capital? At the heart of this debate lies the role of the labor market —both in central bankers' understanding of the inflationary process and in the transmission mechanism of monetary policy. With its explicit dual mandate, the Federal Reserve is a pivotal test case. We measure Fed policymakers' understanding of the labor market as a driver of inflation via LLM-assisted text classification methods, applied to the complete corpus of Fed communications during the period 1978-2019. We document a robust 'rhetorical Phillips curve'—a negative relationship between the rate of unemployment and the salience of labor as a driver of inflation. Going beyond rhetorics, our results show that policymakers' fear is *genuine*—it is more pronounced in internal deliberations of the Federal Open Markets Committee than in public speeches. Lastly, the Fed's fear is *political*—in line with the established finding of a strongly positive effect of Democratic control of the presidency on employment, the Fed is more fearful of full employment when under Democratic presidents.

Keywords: Delegation, distribution, central banks, monetary policy, LLMs.

1. Introduction

Although the United States never built a full employment welfare state, the macroeconomic policy regime it established in the aftermath of World War II put considerable emphasis on the goal of full employment. The centrality of employment is reflected in the dual mandate of the Federal Reserve, which puts monetary policy in the service of both price stability and maximum employment. After a peak during the Volcker disinflation of the early 1980s, a long period of relative macroeconomic stability—the so-called Great Moderation—somewhat blunted the political salience of labor market performance. However, following the global financial crisis, in the course of which unemployment peaked at 10 percent, the issue re-emerged as a central theme in economic policy debates. These debates were fueled by growing income inequality (Piketty & Saez, 2003) and what has come to be called the "China shock"—deindustrialization driven by import competition (Autor, 2015; Autor et al., 2013). Salience has remained high in the context of an extreme spike in unemployment during the Covid-19 pandemic and amidst concerns over the labor market impact of automation and artificial intelligence (Acemoglu & Restrepo, 2022).

Few variables are as decisive for employment outcomes as the Fed's understanding of its monetary policy mandate. Curiously—from a democratic theory perspective—this understanding lies within the central bank's own discretion (Downey, 2024). After the Employment Act (1946) established full employment as a key goal for economic policy, the Humphrey–Hawkins Full Employment Act (1978) instructed the Federal Reserve to pursue both price stability *and* maximum employment (Binder & Spindel, 2017). Humphrey–Hawkins was signed into law a mere twelve months before Fed chairman Paul Volcker, freshly appointed, began administering his eponymous shock. By the end of 1982, unemployment had risen to a post-Great Depression high of 10.8 percent, and the Fed had re-arranged its two goals in a hierarchical fashion, giving de-facto precedence to price stability (Goutsmedt, 2022; Krippner, 2011). In practice, the Fed followed the majority of central banks in embracing inflation targeting (Kaya, 2022). It was in the post-2008 context of heightened political attention on employment that the Fed eventually revised its understanding of "maximum employment", based on its newfound "appreciation for the benefits of a strong labor market" (Powell, 2020, p. 10).

In light of this emphasis on price stability over full employment, critics have described central bank independence as a mechanism to constrain the power of workers and organized labor, and thus to entrench the trend towards increasing income and wealth inequality. Specifically, the Fed has been accused of failing its Congressional mandate by letting its monetary policy be guided by a "baseless fear of full employment" (Galbraith et al., 2007). In spite of a

burgeoning literature on the "politics of central banks" (Adolph, 2018), this accusation has never been subjected to empirical testing. Historical-institutionalist scholarship has shown that the Fed possesses significant discretionary power in both monetary and financial policy (Binder & Spindel, 2017; Downey, 2024; Jacobs & King, 2016; Moschella, 2024). A quantitative strand has shed light on the distributional consequences of central bank independence (Aklin & Kern, 2021), ideological alignment among regional bank presidents (Ainsley, 2021), or the Fed's reputation in Congress (Bellodi, 2023; Ferrara et al., 2021). Still, these major advances notwithstanding, the material question of how the Fed adjudicates between the goals of price stability and maximum employment has remained unanswered.

We therefore ask if—and under what political configurations—the Fed has shown fear of full employment. Specifically, we formulate three hypotheses, of increasing stringency. In its base form, fear of full employment requires that Fed policymakers place greater emphasis on labor as a driver of inflation at lower levels of unemployment, even without signals of accelerating wage growth. In a stronger form, fear of full employment is not merely a rhetorical device but genuinely held: policymakers place greater emphasis on labor as a driver of inflation in their private, decision-relevant deliberations in the Federal Open Markets Committee (FOMC) than in their public speeches. In its partisan form, the Fed's fearfulness reflects Democrats' relatively greater commitment to full employment: policymakers' emphasis on labor as a driver of inflation is greater under Democratic presidents than under Republican presidents.

To test these hypotheses, we deploy large language model (LLM) assisted text classification methods on the complete corpus of Federal Reserve communications during the period 1978-2019. This corpus includes both minutes and transcripts of the internal deliberations of the FOMC and external, public-facing speeches. We operationalize fear of full employment by constructing an indicator of the salience of labor as a driver of inflation across all policymaker pronouncements, and study the associations between relevant macroeconomic, institutional, and political variables and this metric. We find robust empirical support for all three hypotheses.

The paper proceeds as follows. Section 2 situates the concept of full employment in the broader history of post-war economic policy in the United States. Showing how the notion of a trade-off between inflation and employment became embedded in monetary policymaking, Section 3 specifies the hypotheses regarding the Fed's fear of full employment. Section 4 describes the data, coding strategy, and use of LLMs. Section 5 presents the empirical strategy and discusses the regression results.

2. The politics of full employment

The economic policy goal of full employment has had a remarkable career over the course of the past century. The experiences of the Great Depression and of the employment boom of the early 1940s appeared to validate Keynesian economics. By 1944, at the peak of World War II, unemployment in the UK and the United States had fallen to 0.6 per cent and 1.2 per cent, respectively. Full employment was at the core of a newly formed Keynesian policy consensus (Beveridge, 1944). In the context of the U.S. government reducing its massive wartime footprint in the economy, and of millions of soldiers returning home and re-joining the civilian labor force, Congress passed the 1946 Employment Act. Originally introduced as the Full Employment Bill, it had President Truman's support but was significantly watered down in Congress. Manufacturing and farming associations, in particular, lobbied the House and found a receptive audience in a coalition of conservative Southern Democrats and Republicans (Bailey, 1950; Binder & Spindel, 2017, p. 144-150).[1]

Despite this initial setback, full employment was central to US politics during the postwar decades. Public opinion surveys consistently reported strong support for full employment, and the Democratic party routinely included a full-employment plank in its national platform (Weir, 1987, p. 377). The partisan nature of full employment has its roots in the diverging interests of capital and labor. As most famously articulated by Kalecki (1943, p. 324), the question of why business leaders tend to oppose Keynesian policies geared towards full employment is "not easy to explain." After all, higher economic output increases business profits. However, capitalists value the disciplining force unemployment exercises on organized labor, and fear that "under a regime of permanent full employment, 'the sack' would cease to play its role as a disciplinary measure" (Kalecki, 1943, p. 326). In describing the regularity with which successful full employment policies are squashed by business opposition, Kalecki coined the concept of a "political business cycle."

After falling throughout the 1960s, unemployment shot up during the first half of the 1970s. For the first time, the US economy experienced stagflation. Inflationary pressure escalated with the oil crisis 1973-74, during which the price of crude oil more than doubled within weeks. It was in this context of unprecedented stagflation, that the Humphrey–Hawkins Full Employment Act of 1978 "rejected the passive and reactive Fed" of the past and instead directed it "to foster and sustain economic conditions conducive to maximum employment,

[1] At the time of the Employment Act's passing, full employment was a prominent topic in political science. The December 1945 issue of the *American Political Science Review* featured a symposium on "Maintaining High-Level Production and Employment" with several contributions on the topic. However, this moment was fleeting. Indeed, just like it disappeared from the legislation's name, after 1945, the phrase "full employment" never re-appeared in an APSR or AJPS research article title.

production, and purchasing power of the dollar" (Binder & Spindel, 2017, p. 150). However, the same stagflation crisis that paved the road for the original bill in 1974 also paved the road for Paul Volcker to become chairman of the Fed a mere 12 months after President Carter had signed the Humphrey–Hawkins Act into law.

Volcker's immediate embrace of a radical disinflation policy permanently relegated full employment to second place in the hierarchy of the Fed's dual mandate (Kaya et al., 2019). Volcker's Fed raised rates to create a severe recession and a surge in unemployment, a strategy recently qualified by scholars as "Keynesianism in reverse" (Best, 2020; Binder, 2021) . Based on a cost-push theory of inflation, according to which "labor accounts for the bulk of all costs" underpinning "the momentum of the inflationary process" (Volcker, 1981), Volcker was intent on breaking the structural power of organized labor in order to achieve permanent price stability (Greider, 1987, p. 431). In Volcker's view, President Reagan made an "important but little-recognized contribution to the fight against inflation" when he crushed the air traffic controller strike, thus sending "a powerful psychological message that there would be limits on wage demands" (Volcker & Harper, 2018, p. 113). In short, there can be little doubt that Volcker's anti-inflation strategy was to fundamentally reshape the U.S. economy to the disadvantage of labor. This strategy succeeded–the Volcker shock "all but erased" full employment as a major political issue in the United States (Weir, 1987, p. 377).

As full employment receded as a policy priority, the labor market retained a central place in central bankers' minds, albeit in negative relief. Full employment was no longer seen as a goal in itself, but instead as a potential threat to price stability. Underpinning this view is the "Phillips curve", which in its original formulations posited a negative relationship between unemployment and wage growth (Phillips, 1958), and subsequently the rates of unemployment and inflation (Samuelson & Solow, 1960). The microeconomic foundations were added by Milton Friedman (1968) and Edmund Phelps (1967, 1968) who, in the process, developed the notion of a "natural" rate of unemployment at which inflation remains anchored, yet below which wage demands escalate and wage-price spirals become likely. In its "modern", New Keynesian form, the Phillips curve incorporates forward-looking expectations and stands for a trade-off between the amount of "slack" in the economy and inflation. From this perspective, inflation is the product of "excess" aggregate demand relative to productive capacity (Roberts, 1995).

This New Keynesian consensus provided a theoretical bulwark against full employment Keynesianism—"an expectations-based argument against attempts to lower unemployment by demand management" (Forder, 2010, p. 330). Some critics went further, accusing monetary policymakers of a "baseless fear of full employment" (Galbraith et al., 2007). Did the Fed

conduct its monetary policy on the basis of such a fear? In other words, did it systematically fail to "appreciate the benefits of a strong labor market"?

3. Full employment, monetary policy, and the Fed

Assessing this question empirically is fraught with difficulty. Central bankers with a democratic mandate to pursue maximum employment are in no position to state their fear of full employment plainly. In a moment of unusual candor at the peak of the dotcom boom, Fed chairman Greenspan (1999) warned of the inflationary consequences of the "steadily depleting the pool of available workers", which could not "continue without eventually putting increasing pressure on labor markets and on costs." Once nominal wages start growing faster than labor productivity, he argued, "prices inevitably will [...] begin to accelerate." Greenspan's statement points to a clear trade-off in his mind between full employment and inflation. The implication that disinflation requires higher unemployment was stated by Larry Summers in ..., when he claimed that "we need five years of unemployment above 5 percent to contain inflation" (Blyth & Fraccaroli, 2025). Statements such as these are highly significant because they point to the existence of fear of unemployment among economists and central bankers.

It was only in the aftermath of the 2008 financial crisis that the post-Volcker prioritization of inflation over full employment came under increased scrutiny. After decades during which the concept of "full employment" had been largely absent in FOMC discussions, policymakers began to refer to it much more frequently during the early 2010s. Fed policymakers also began to question the concept of the natural rate of unemployment, which culminated in the Fed conducting an official strategy review process in 2019-2020 (Arbogast et al., 2024). As chairman Jerome Powell noted upon the conclusion of the strategy review, the Fed's revised understanding of the maximum employment mandate reflected its "appreciation for the benefits of a strong labor market, particularly for many in low- and moderate-income communities" (Powell, 2020, p. 10). For the first time since Volcker, the Fed set its fear of full employment aside and allowed the labor market to run hot.

In sum, this survey of the literature points to the ability of the Fed to exercise significant discretionary power in the understanding and pursuit of its maximum employment mandate. It also, however, points to a methodological challenge: How to measure the views of policymakers' who are keen to avoid overly candid statements of their preferences for employment? To overcome this challenge, we develop an indicator based on central bankers' favorite way of discussing employment—namely, as one driver of inflation among others. Relating this indicator to real-world macroeconomic variables, such as the unemployment rate or wage growth, allows us to construct a proxy indicator for central bankers' fear of full employment.

To test the fear of full employment proposition, we formulate three hypotheses. At a minimum, demonstrating fear of full employment at the Fed requires that policymakers react to increases in employment rates *even in the absence of inflationary wage growth*, which in the Phillips curve logic transmits labor market pressure into generalized consumer price inflation. Thus, *we expect that Fed policymakers place greater emphasis on labor as a driver of inflation at lower levels of unemployment, regardless of the level of wage growth (**H1**).*

Protective of the fragile institution of central bank independence, central bankers have been shown to be engaged in sophisticated reputation management strategies (Moschella, 2024; Romelli, 2022; 2024). Notwithstanding frequent references to the importance of transparency in central bank rhetoric, discrepancies between central bankers' beliefs and what they say publicly in their speeches have been shown to be ubiquitous (Braun, 2016; Diessner, 2023; Goutsmedt & Fontan, 2024). This is especially true when central banks embrace a monetary policy stance that has negative consequences for employment (Best, 2020). Therefore, policymakers *not* expressing their fear of full employment candidly in public would not contradict the fear hypothesis. To the contrary—greater candor in private settings would indicate that the sentiment is genuine. Thus, our second hypothesis is that *policymakers' fear of full employment is genuine—they place greater emphasis on labor as a driver of inflation in their private, decision-relevant deliberations in the FOMC than in their public speeches* (**H2**).

Full employment is a matter of class politics, and thus a partisan issue (Kalecki, 1943). Left-wing parties, when in government, place a higher premium on employment than conservative parties. During the postwar decades, left-wing governments were associated with high employment-high inflation outcomes, whereas right-wing governments presided over high-unemployment-low inflation outcomes (Hibbs, 1977). This employment policy partisan split has had profound economic consequences for workers. In the US, middle-class families have seen their real incomes increase "more than twice as fast under Democratic presidents", while for working poor families that growth rate has been "ten times as fast under Democrats as under Republicans" (Bartels, 2016, p. 37). Relatedly, partisan control of state government has been shown to have significant employment effects along racial lines, with annual hours worked by black workers increasing relative to white workers under Democratic governors (Beland, 2015). Importantly, the partisan split extends beyond employment to GDP growth. Regardless of how economic performance is measured, US economic performance has been better under Democratic presidents than under Republican presidents, with a particularly large gap for real GDP growth (Blinder & Watson, 2016).

Given the greater dedication of left-wing parties to the goal of full employment, it is rational for central bankers to be *more fearful* of full employment under left-wing governments.

Crucially, this should hold *regardless* of policymakers' own partisan affiliation or preferences, which have been shown to exert a significant influence on monetary policy decisions (Pagliuca, 2025). Our argument here follows the logic in Cusack (2001), who has shown monetary-fiscal coordination to be less likely under left-wing governments; and in Clark & Arel-Bundock (2012), who have shown Fed policymakers to be "conditional inflation hawks" whose monetary policy stance is more hawkish under Democratic than under Republican presidents. Thus, we expect fear of full employment to have a partisan dimension—*policymakers' emphasis on labor as a driver of inflation is greater under Democratic presidents than under Republican presidents* (**H3**).

4. Using LLMs to detect causal claims about inflation

To test our hypotheses, we need to detect when Fed officials discuss labor-related issues as a cause of inflation. First, we collect the transcriptions of all public speeches by Fed officials as well as internal deliberations in the FOMC, and extract excerpts in which inflation is discussed. Then, we create a codebook to categorize the different drivers of inflation in Fed discourse, and use it to label a subsample of inflation excerpts. Last, we benchmark several LLMs for our classification task, and use the most capable to systematically label inflation drivers across our whole corpus of Fed discourse.

4.1. Public and private discussions of inflation

The FOMC, the Federal Reserve's monetary policy decision-making body, consists of twelve members: the seven members of the Federal Reserve Board of Governors, the president of the Federal Reserve Bank of New York, and four of the remaining eleven Reserve Bank presidents serving one-year rotating terms, while the rest participate in meetings but do not vote. The FOMC meets at least eight times a year to review economic conditions and determine monetary policy stance, with additional conference calls happening in times of crisis. FOMC transcripts are available from March 1976 to December 2019 (as of 2025), and are released with a delay of five years.[2] There are 477 transcripts in total (367 regular meetings and 110 conference calls). After downloading these documents in a *PDF* format, we leverage natural language processing tools to parse these transcripts into individual speakers' interventions, retaining only those that were delivered by Board members and Reserve Bank presidents, excluding staff presentations. This yields 123,039 interventions averaging 96 words, leading to 12M words total.

[2]FOMC transcripts began to be made public in 1993.

In addition to these private deliberations, Fed officials give public speeches, in which they explain their decisions in order to guide the expectations of markets and citizens. We collect the text of these speeches from two sources. First, we use the Central Bank Speeches (CBS) database by Campiglio et al. (2025), focusing on the Federal Reserve. For this institution, CBS aggregates Fed speeches from 1986 onward from three sources: the Bank of International Settlements repository, Federal Reserve websites (Board and regional Feds), and the St. Louis Fed FRASER archives. Second, we extend coverage to 1970-1986 through additional FRASER scraping, matching the existing data format. Limiting our sample to Board members and regional presidents yields 8,253 speeches averaging 3,010 words, or approximately 24M words.

We then need to separate text about inflation and price-related issues from other irrelevant discussions. Indeed, speeches may only address inflation in subsections of the speech, or even avoid the topic entirely. To focus on relevant discussion only, we split our collection of interventions and speeches into sentences,[3] and then implement a keyword-based search to identify sentences that discussion inflation or price dynamics using a list of 16 common expressions including, for example, *inflation*, *pce*, and a range of bigrams including *price* (a full list of dictionary terms and the methodology used to identify them is available in Appendix A.1). This search yields 129,457 relevant sentences containing one or several of our inflation-related keywords.

Since discussions of inflation drivers often extend beyond individual sentences, we adopt a classic method in computational text analysis, aggregating relevant sentences with the sentence before and after to create sentence triplets, providing more context.[4] To prevent double-coding when sentences appear as context in multiple excerpts, we then aggregate consecutive and near-consecutive relevant sentences into cohesive, longer excerpts. This methodology yields 60,150 inflation excerpts containing our 129,457 inflation-related sentences. The majority of excerpts are 3 sentences long, with a few being shorter (when FOMC interventions contain fewer than 3 sentences), and some being longer (due to our aggregation method). Since our dictionary of inflation-related keywords was extensive, to minimize false negatives at the cost of potential false positives, we are confident these excerpts capture the vast majority of any inflation-related discussions in either private deliberations or public speeches.

[3] We use the python package *spacy* and its largest parsing-trained model *en_core_web_lg* to do this.

[4] We tested larger windows (two or three sentences on either side) but found they increased coding volume without capturing additional relevant information. Subsequent relevant sentences typically contained our target keywords and were already identified.

4.2. LLM classification of inflation drivers

Our goal with these inflation excerpts is to detect when Fed officials discuss labor market dynamics as causing price increases. A naive approach would simply count mentions of labor-related keywords within these inflation excerpts. Yet keyword co-occurrence cannot distinguish between fundamentally different semantic relationships: "Rising wages are driving inflation" attributes causality to labor, while "Inflation is eroding workers' wages" discusses the consequences of price increases for labor. Words do not exist in isolation: their meaning emerges from the web of relationships they form with other concepts in discourse, something that legacy natural language processing techniques based on bag of words approaches are largely enable to detect, but that large language models may be capable of identifying (Garg & Fetzer, 2025).

Yet, using large language models for text annotation tasks requires substantial human work to ensure that model outputs are indeed satisfactory approximations of what researchers aim to capture. Following best practices in the literature, we implement a robust pipeline including (1) exploring a *test sample* of several hundred inflation excerpts to design a codebook, (2) classifying a larger *validation sample* to benchmark model performances against a set of human labels, and (3) selecting an appropriate model and exploring several avenues to further enhance performance (Cova & Schmitz, 2024; Grimmer et al., 2022; Törnberg, 2024).

Accurately capturing central bankers' causal claims about inflation requires a clearly defined categorization scheme. To make the task as objective, replicable and robust as possible, we first build a codebook of inflation drivers. Since labor-related factors represent only one potential explanation for inflation among many, we construct a categorization capable of capturing the full range of causal claims central bankers advance to account for price increases.[5] Based on the literature on the sources of inflation in US history, we build a preliminary codebook that includes not only labor, but also other prevalent explanations of price increases, such as the government's fiscal stance, the exchange rate, or energy and commodity prices (Alvarez et al., 2022; Blinder, 2022; Blyth & Fraccaroli, 2025; Drechsler et al., 2022; Goutsmedt, 2021; Judge, 2023). We then progressively drew a *test sample*, including 275 inflation excerpts in total, which was coded independently by each of the three authors in order to develop this codebook. Through this process, results were compared and categories refined iteratively until classifications appeared consistent, suggesting a common understanding of the task and a proper delimitation of the labels. Our final categorization includes six main inflation

[5]This has two benefits: putting labor-based explanations of inflation in perspective, and enabling future work on the other drivers of inflation in Fed discourse.

drivers (including "Labor") and a residual category aimed at capturing other, less frequent explanations for inflation (see Appendix A.1 for the full codebook).[6]

We then move to selecting a model. We drew a larger set of 800 new randomly drawn inflation excerpts (400 from speeches, 400 from FOMC meetings) which were again coded manually to create a *validatation sample* to serve as a benchmark to assess the performance of different LLMs. Based on our test sample, we accurately anticipated that only approximtely a quarter of inflation excerpts would contain causal claims, and with 800 excerpts we achieved the guideline provided by Törnberg (2024) that LLM performance should be validated on a minimum of 20-30 examples per category. For our "Labor" code, we had 71 positively coded excerpts in this sample. Each excerpt was coded independently by two different authors, with disagreements resolved through consensus with the third author. Disagreements are not very frequent: the inter-coder reliability score (Krippendorff's Alpha) for Labor was 0.81.

Several families of large language models (LLMs) could be leveraged for our classification task (Do et al., 2022). We select generative LLMs over transformer-based classifiers like BERT for three advantages: they require no manually-coded training data[7], they can provide rationales for classifications, and they are able to handle complex instructions via natural language prompts. We test twelve different generative LLMs from different providers, including closed and open source models (e.g. DeepSeek-V3 vs Claude Sonnet 3.5), as well as reasoning or legacy models (e.g. o4-mini vs GPT4o). Closed-source models may offer superior performance but present reproducibility challenges and privacy concerns, while open-source alternatives improve transparency at the cost of potential performance limitations. Reasoning models employ multi-step chain-of-thought decomposition, which can improve performance on complex interpretive tasks but increases token generation costs and may introduce "overthinking" for classifying text.

All models are given the same prompt (a set of instructions in natural language, similar to the codebook that was designed for human coders), and their output is then compared to our human labels. The performance of each model is measured for each binary task (e.g. is the excerpt discussing labor as driving inflation?) with *F1 scores* (the harmonic mean of *precision*

[6]For a similar typology of inflation narratives, see Fraccaroli et al. (2025), who highlight public spending, wage increases above productivity growth, supply side disruptions, and corporate profit margin expansion.

[7]This advantage is especially relevant for tasks *unbalanced* tasks like ours: since our labels appear in less than 10% of excerpts, a training dataset big enough to provide enough positive examples to the model on each of the categories would require the coding of a prohibitively large amount of texts by human coders.

[8]Precision measures how often the model is correct when it identifies a category, while recall measures how many true instances it successfully captures. The F1 score balances these metrics, penalizing models that excel at one dimension while failing at the other. This is particularly important for our classification tasks,

and *recall*) rather than a simple accuracy metric.[8] Figure 1 displays the F1 scores across all twelve models.

Figure 1: F1 scores

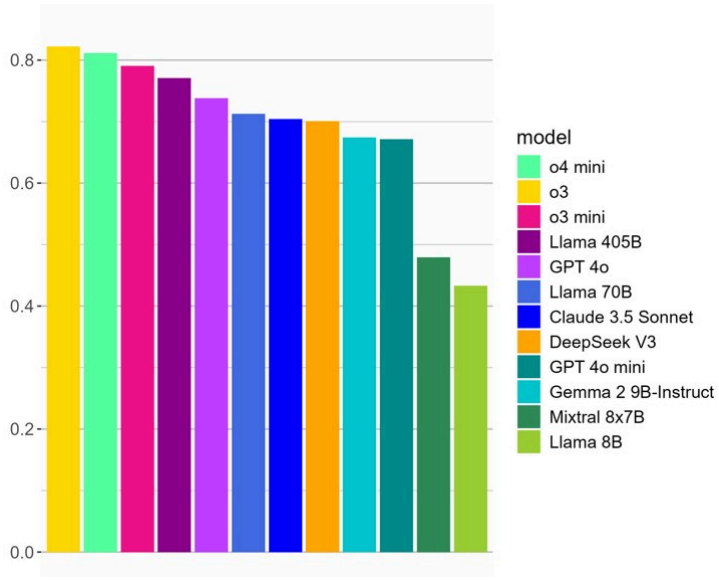

Note: Weighted averages across the six categories in our codebook.

Unfortunately, open source models are largely outperformed by closed competitors, and do not systematically score above 0.7, which leads us to select *o4-mini*. We decide to pick this model over *o3* because it performs slightly better on our main category of interest (*Labor*), while being four times cheaper in inference costs. We provide detailed performance metrics for this model as well as o4-mini and the two best open-source models (Llama 3.1 405B from Meta and DeepSeek-V3) in Appendix A.1. Before sending our full corpus to OpenAI's API, we conduct a last step of performance validation and explore possibilities to further improve o4-mini's performance. This included "prompt-tuning", testing a few changes in the way we instruct the model to see if performance can be improved, as well as trying few shot promoting (where example excerpts are included in the prompt) and both mixture-of-experts majority voting using three of our highest performing models (o3-mini, GPT4o, and Llama-3.1-405B) as well as single-model majority voting. However, none of these measures yielded a significant performance increase, suggesting that gains from prompt tuning may be lower with the latest generative LLMs. As a final step, we also used the *promptstability* package from Barrie et al. (2025) to test the stability of our task across multiple runs of our model for our validation sample, with intra-prompt stability scores reported in Appendix A.1. The final instructions provided to the model are also available in Appendix A.1.

in which the prevalence of each individual inflation driver to appear is low: a model could achieve high accuracy by never predicting a category, making F1 scores more meaningful than simple accuracy.

The output of the classification by *o4-mini* includes 23,259 excerpts with at least one inflation driver, out of the 60,150 inflation related excerpts in our corpus. *Labor* appears as a cause of inflation in 6,172 excerpts (about 10%) of inflation excerpts, making *Labor* the main explanation of inflation in our Fed corpus. Disaggregating this further reveals significant temporal variability, as displayed in Figure 2 which plots both the absolute quarterly count and relative salience (the count divided by total words spoken) of excerpts mentioning Labor as a driver of inflation.

Figure 2: Labor time series

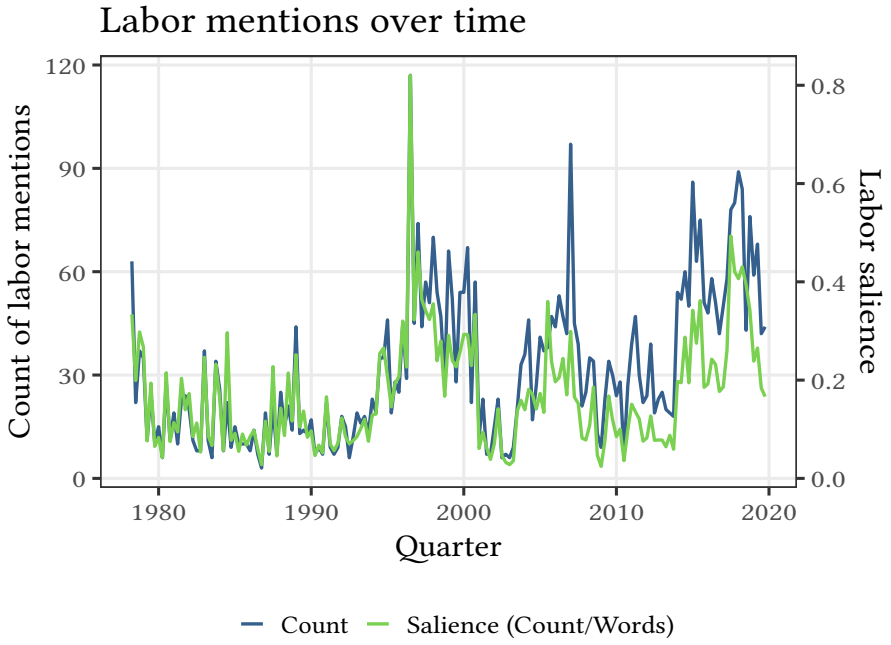

Note: Relative importance of inflation drivers (left) and labor mentions over time (right)

5. Evidence for fear of full employment

We can examine several bivariate plots to investigate if there may be an intuitive case for a fear of full employment using our newly constructed LLM dataset. While we can only formally test this relationship through the statistical analysis in subsequent sections, these plots provide a compelling initial case on the relationship between unemployment and the volume of Fed officials' discussion of labor market dynamics as inflationary. Figure 3 plots our aggregate quarterly count of mentions of Labor as a driver of inflation, divided by total quarterly word count in our corpus, against unemployment figures from the Bureau of Labor Statistics.[9] Already, the existence of at least a rhetorical Phillips curve seems plausible: there is

[9]As explained in Section 5.1, we lag many of our macroeconomic variables by one quarter to account for release dates.

a clear inverse relationship, with an apparent elbow around 6%, below which Labor mentions increase substantially.

Figure 3: A rhetorical Phillips curve?

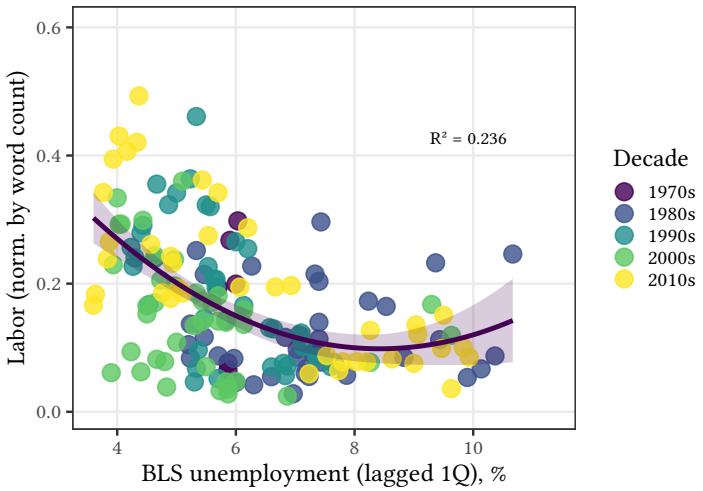

Note: One Labor outlier removed

Figure 4 creates the same plots, but this time disaggregated by forum. Recall that our second hypothesis posits that if the fear of full employment is genuine, we would expect a more intense pattern of communication in internal FOMC policy deliberations, compared to publicly facing speeches. Again, there seems to be plausible evidence that this might be the case: mentions of Labor as an inflation driver in the plot on the left, for FOMC meetings, increase much more steeply at low levels of unemployment as compared to the plot on the right, for public speeches, which looks relatively flat in comparison.

Figure 4: Labor salience by forum

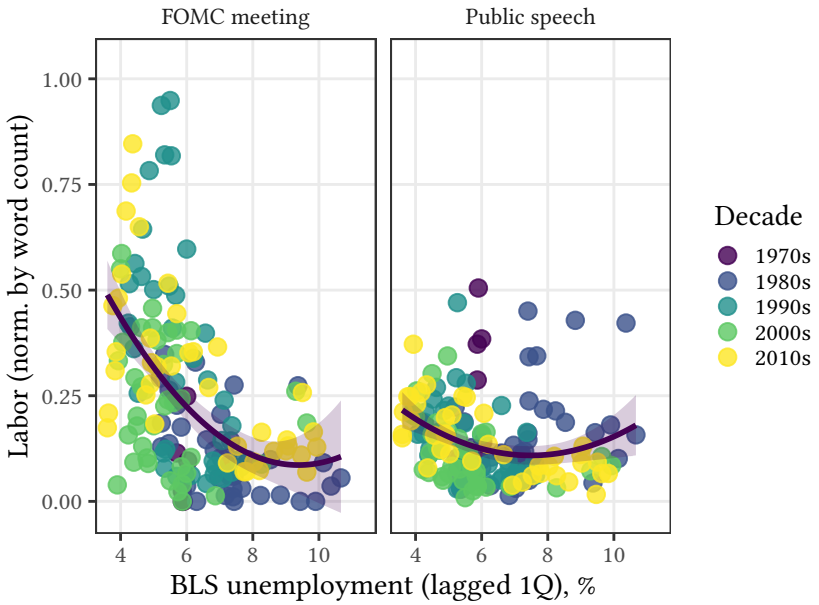

Note: One Labor outlier removed

5.1. Empirical strategy

While these plots provide compelling descriptive evidence, we now move to formally test our hypotheses. We estimate three different sets of models to correspond with the three hypotheses presented in Section 3. Our analysis spans from 1978-Q2 until 2019-Q4.[10] We begin with a baseline specification that corresponds with H1: do Fed officials exhibit a fear of full employment? More specifically, we examine this empirically by studying the effect of changes in the unemployment rate and wage growth on *Labor* mentions in Fed communications. We estimate the following model:

$$\text{Labor}_t = \beta_0 + \beta_1 U_{t-1} + \beta_2 U_{t-1}^2 + \beta_3 N_t + \Omega X_{t-1} + \alpha_c + \varepsilon_t \tag{1}$$

Our dependent variable, Labor_t, is the count of Labor-coded excerpts aggregated for a given quarter t in all Fed communications, as presented in Section 4. Summary statistics for our dependent variable are presented in Table 1.

Table 1: Dependent variable summary statistics

Variable	Mean	SD	Min	Max	Variance
Labor	31.4	21.87	3	117	478.45

[10]We begin our analysis in 1978-Q2 instead of 1978-Q1, even though the Fed's dual mandate had already been signed into law as of Q1, because 1978-Q1 is the last quarter where Arthur Burns was Chair before the end of his tenure. Given our examination of speech dynamics under different Chair periods, we begin our analysis with the tenure of G. William Miller in 1978-Q2, the first full quarter where he was Chair.

The main independent variable, U_{t-1}, is the unemployment rate provided by the Bureau of Labor Statistics (BLS), lagged by one quarter to reflect the delay in release dates (see Appendix A.2 for further details on all data sources). U_{t-1} is centred on its mean.[11] We also test all of our specifications using projected unemployment rates for a given quarter t provided contemporaneously to Fed policymakers by staff in Tealbooks (formerly Greenbooks) ahead of FOMC meetings in the Appendix. We use Tealbook data to reflect the institutional outlook on unemployment rates that should theoretically be influential on the views of Fed policymakers and their speech. We furthermore test the significance of a quadratic term for unemployment, U_{t-1}^2, given the focus of existing economics literature on the presence of a NAIRU, the unemployment rate below which inflation is expected to rapidly increase. If Fed officials do indeed have a fear of full employment, we would expect the coefficient on U_{t-1} to be negative and significant (as unemployment falls, *Labor* mentions should increase). If this effect accelerates as unemployment falls, we would expect the coefficient on U_{t-1}^2 to be positive and significant. To capture the possibility of shifting structural conditions in labor markets over time, we also test specifications that instead substitute a metric for the *Unemployment gap*, which we explain in further detail in Section 5.2.

We also estimate specifications that include further macroeconomic covariates, X_{t-1}, which may also influence the extent to which Fed officials discuss Labor as a driver of inflation. First, and crucially, we include specifications that account for the effect of wage growth. If the mechanism by which Fed officials are concerned that falling unemployment may increase inflation is wage growth, we would expect coefficients on wage growth variables to be positive and significant. We test two specifications that include variables for *Nominal wage growth* and *Real wage growth*, respectively. In the specifications in the main text, we use Average Hourly Earnings (AHE) data provided by the BLS to create our wage growth variables, which we lag by one quarter.[12] In the Appendix we also test wage growth variables using Median Weekly Earnings (MWE).[13] Second, we also include *Inflation* as a covariate, given that we may expect more frequent references to Labor as a driver of inflation when inflation rates are higher.[14] Finally, we also include *GDP growth*, as Fed officials may be more concerned about Labor as a driver of inflation in periods where output growth is high. All macroeconomic covariates are

[11] We centre macroeconomic variables on their means to reduce multicollinearity issues as we progressively add quadratic and interaction terms to our model specification.

[12] AHE data is collected monthly. For *Nominal wage growth*, we average wage data across the quarter, and then take the percentage change from four quarters ago. For *Real wage growth*, we deflate nominal wages using the CPI to convert them into constant 1982-84 dollars, and then follow the same methodology.

[13] MWE data is also provided by the BLS, but is only available beginning in 1979, so in these specifications we drop the first year from our dataset and begin our analysis in 1979-Q1.

[14] However, we do not include a control for *Inflation* in specifications that include *Nominal wage growth*.

also centred on their means. The models estimated in the main text use data from the Bureau of Economic Analysis (BEA) for *Inflation* and *GDP growth* (lagged by one quarter). We also estimate models using Tealbook projections for these variables in the Appendix. The vector Ω_{t-1} contains the coefficient estimates for these covariates. See Appendix A.2 for further details on macroeconomic variables and their sources.

Finally, we also include a control for the volume of total communication, N_t, which is the logged word count for our entire corpus of Fed communication in a given quarter, given that the number of mentions of Labor as a driver of inflation in a quarter is likely to correspond with the overall volume of communication. We also include dummies for Fed Chairs, α_t, in each of our specifications.[15] This allows us to control both for structural breaks in the relationships between these variables over our entire sample, as well as the hypothesis that different Chairs are influential in the shape that communication at the Fed takes in aggregate during their tenures. β_0 is a constant, and ε_t is the error term.

Given that Labor_t is a count variable and exhibits overdispersion (see Table 1), we estimate Equation (1) using a negative binomial regression, as well as all subsequent model iterations included in the main text. However, we also estimate our specification using several different models (including Poisson and quasi-Poisson) in the Appendix.

5.2. A fear of full employment

Figure 5 presents the results for the model specifications described in Section 5.1. All models include *Unemployment* variables, with model specifications progressively including other key macroeconomic variables including *Inflation*, *Wage growth*, and *GDP growth*.[16] All models include controls for total volume of words, as well as Chair fixed effects. Full model results are available in Table 4.

[15]There are six Fed Chairs during our sample period: G. William Miller, Paul Volcker, Alan Greenspan, Ben Bernanke, Janet Yellen, and Jerome Powell.

[16]The model in column 3 contains only nominal wage growth, which is constituted by effects from both inflation and real wage growth, whereas the model in column 4 includes variables for real wage growth and inflation separately.

Figure 5: Model coefficients for H1, 1978-2019

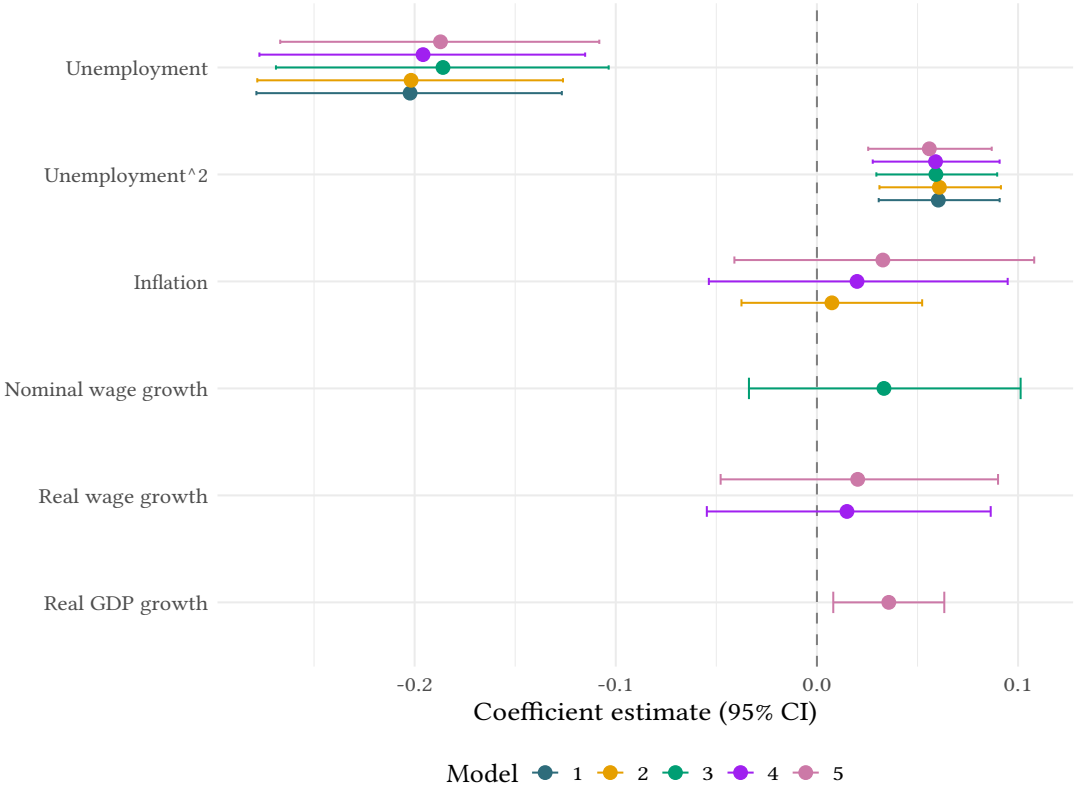

Both *Unemployment* variables are highly consistent in sign and significance across all specifications. The sign on *Unemployment* is negative, indicating that mentions of Labor as a driver of inflation increase in Fed speech when *Unemployment* falls. Furthermore, the sign on the quadratic *Unemployment* variable is positive, indeed indicating accelerating attention to Labor at lower values of *Unemployment*. For example, the model estimates that a fall in unemployment from its mean in our sample (6.22%) to 1% below its mean is associated with an increase from approximately 28 mentions of Labor as a driver of inflation to 35 (an increase of 25%). A further drop in unemployment from 1 to 2% below mean results in a further increase to 50 Labor mentions (an increase of 43%).

Notably, none of the wage growth variables in any of the specifications achieve significance. This is particularly surprising given that wage growth is ostensibly the transmission channel by which low levels of unemployment lead to rising inflation. However, when looking at the simple bivariate plots of real and nominal wages against *Labor* salience in Figure 6 (where we would expect to see a positive relationship if wage growth were significant), this result seems entirely plausible. Furthermore, including these and additional covariates (inflation, GDP growth) does not diminish the significance of the *Unemployment* variables in our model.[17]

[17]While less directly relevant for our hypotheses, the only macroeconomic control for which we find consistent significance is GDP growth. This may reflect some level of Okun's Law at play, demonstrating a

Taken together, these results provide support for a fear of full employment at the Fed. Not only is the effect of *Unemployment* consistent across all specifications, but it is a better predictor of the volume of speech on Labor as a driver of inflation than other key macroeconomic indicators. While *Labor* speech shows a robust relationship with unemployment rates, wage growth has no significant effect on Fed officials' concerns.

Figure 6: Wage growth bivariate plots

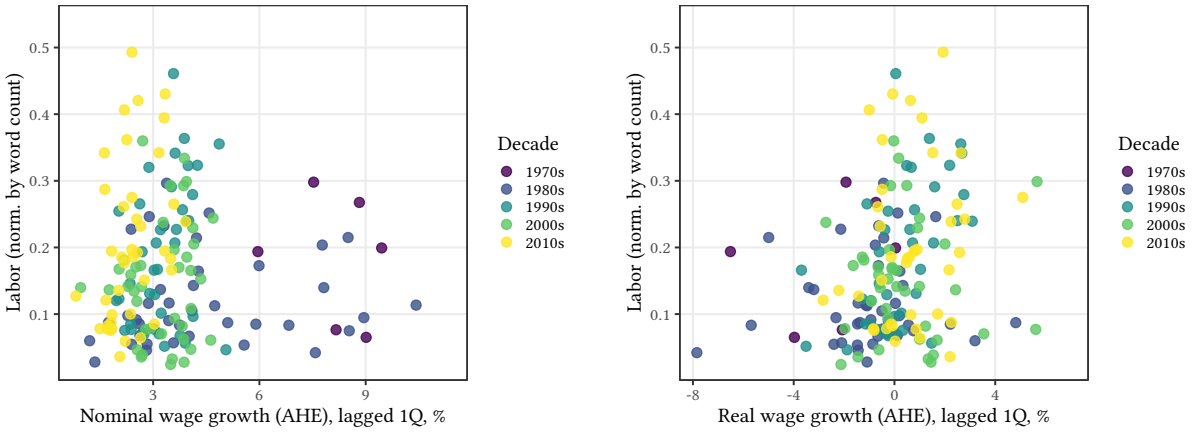

Note: One Labor outlier removed; one real wage growth outlier removed

We conduct several robustness checks on this analysis in Appendix A.3. First, as elaborated in Section 5.1, our results are also robust to other modeling strategies (Table 5). Second, these results are also robust to alternative data sources, including both Tealbook forecasts for macroeconomic variables (Table 7), and an alternative metric of wage growth, the MWE, both of which yield results consistent with those in the main text (Table 8).

We also test robustness by substituting our unemployment variables for the unemployment gap (UGAP) — the difference between the unemployment rate and the Congressional Budget Office's estimated NAIRU for each period (Figure 17). This provides an additional avenue to capture the possibility of structural, supply-side shifts in labor market dynamics over the course of our sample, which may affect the way that Fed officials react to different levels of unemployment over time. This matters because the NAIRU fell steadily over our sample, from approximately 6.2% in 1978 to 4.4% in 2019. A rate of 5%, for instance, represented a tight labor market by 2015 standards but slack conditions by 1982 ones. The UGAP metric thus captures unemployment relative to what policymakers at the time would have considered full employment. However, we find similar results when substituting the UGAP metric in our model specification, providing evidence that our findings are not being driven by the secular decline in the NAIRU (Table 9).

rhetorical link between GDP growth and Fed officials' attention to labor dynamics, albeit in an inflationary rather than growth context.

5.3. Genuine fear of full employment

While a fear of full employment appears compelling from this analysis, we have a further source of variation in our data that we can leverage to test a stronger version of our hypothesis: we can separate our corpus by private FOMC deliberations and publicly delivered speeches. As set forth in H2, if the Fed's fear of full employment is genuine, we might expect a larger effect in FOMC deliberations compared to public speeches, or at least not a weaker effect. We test this by running our regressions from Section 5.2 again, but this time with our dependent variable separated by forum (creating a balanced panel dataset with one observation for FOMC and public speeches in each quarter). We include a dummy for *Forum* (Forum = 1 for FOMC) interacted with *Unemployment* to test whether Fed policymakers' reaction to changes in the unemployment rate is different among public versus private settings. Table 2 presents these results.

Table 2: Results on H2, a genuine fear of full employment, 1978-2019

	1	2	3	4	5
Unemployment	−0.128‡	−0.128‡	−0.114†	−0.125‡	−0.113‡
	(0.042)	(0.043)	(0.045)	(0.044)	(0.043)
Unemployment^2	0.061‡‡	0.061‡‡	0.060‡‡	0.060‡‡	0.057‡
	(0.017)	(0.017)	(0.018)	(0.018)	(0.018)
Forum (FOMC=1)	0.194	0.197	0.205	0.198	0.173
	(0.112)	(0.112)	(0.112)	(0.112)	(0.111)
Unemp*Forum	−0.158‡	−0.158‡	−0.155‡	−0.157‡	−0.164‡
	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)
Unemp^2*Forum	−0.002	−0.003	−0.004	−0.003	−0.003
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
Real wage growth				0.009	0.014
				(0.030)	(0.030)
Nominal wage growth			0.032		
			(0.031)		
Inflation		0.006		0.014	0.025
		(0.019)		(0.032)	(0.032)
Real GDP growth					0.035‡
					(0.013)
Quarterly word count (log)	0.452‡‡	0.454‡‡	0.456‡‡	0.454‡‡	0.436‡‡
	(0.100)	(0.101)	(0.100)	(0.101)	(0.099)
Num.Obs.	332	332	332	332	332
AIC	2342.7	2344.6	2343.7	2346.5	2341.2
BIC	2392.2	2397.9	2396.9	2403.6	2402.1
Chair dummies	✓	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

As before, the coefficients on the *Unemployment* variables, which now represent the effect for public speeches, are significant and maintain the same signs across all specifications: negative on the linear term, reflecting more *Labor* mentions as unemployment falls, and positive on the quadratic term, reflecting an exponential increase in *Labor* mentions at low levels of unemployment. The coefficient on the interaction term *Unemp*Forum* is also negative and significant across all specifications. This provides evidence that Fed officials do indeed react more aggressively to falling unemployment rates in the FOMC than they do in public speeches. While a drop in unemployment rates from 1 to 2% below mean is associated with only five additional Labor mentions in public speeches, it is associated with approximately 12

additional mentions in FOMC deliberations in a given quarter. We again do not find statistical significance for the wage variables included in any of the specifications in Table 2.

We can visualize this difference in effects in Figure 7. While the model predicts a similar volume of discussion of Labor as a driver of inflation in both forums when unemployment is near its mean, as unemployment falls below its mean the rate at which Fed officials begin to discuss Labor (the slope in Figure 7) in private FOMC deliberations accelerates dramatically compared to public speeches. In line with H2, this provides evidence for a *genuine* fear of full employment: while Fed officials may be more cautious in their public communications, they privately express ample concern that labor market dynamics may be inflationary at low levels of unemployment, uncorrelated with wage growth.

Figure 7: Predictions, Model 5

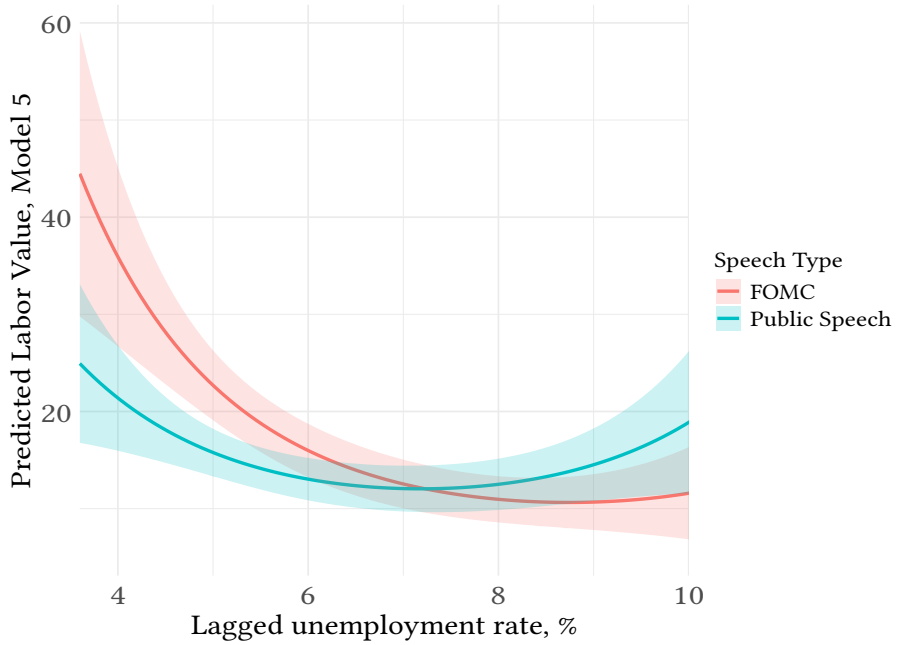

Note: Model predictions calculated while holding other variables at their means and averaging across chairs.

We run several robustness checks on these results in Appendix A.4, again using alternative metrics for our variables including Tealbook data (Appendix A.42) and an alternative wage metric, the MWE (Appendix A.43), and the UGAP (Appendix A.44) and find no change in our substantive results: the *Unemployment* and *Forum* variables remain significant in all and of comparable effect size, and wage variables do not gain significance.

We also probe an alternative explanation for our results. It could be possible that the differences in communication patterns we observe in FOMC meetings compared to public speeches could be driven by systematic differences in speaker composition in these different forums, rather than the effect of the forum itself. For example, if regional Fed presidents are

thought to more closely represent business interests, we might expect that they would be more hawkish on labor than Board members. It could be possible that labor-hawkish regional presidents speak in similar volumes about labor as a driver of inflation in each forum, but they speak relatively less in total volume in public speeches (or vice versa). If this were the case, we would expect (1) a substantively different speaker distribution in each forum, (2) differences in labor salience between groupings of Board members and regional presidents, and (3) the significance of *Forum* to disappear in subgroup regressions. The descriptive plots in Appendix A.45 do not provide strong evidence for (1) or (2), and when we re-run our regression from Section 5.3 separately for the Fed Board and regional presidents as subgroups, *Forum* remains significant in both (also in Appendix A.45).[18]

One other event also bears mention, that it only became known to Fed policymakers in 1993 that FOMC transcripts existed dating back to 1976, that those would be made public, and that going forward transcripts would be made public with a five-year lag. While we would expect this development to make Fed officials *less* likely to discuss Labor as a driver of inflation in FOMC transcripts, bringing them more in alignment with public speeches, we nonetheless tested the inclusion of a dummy for 1993-Q4 in our regressions for Section 5.3. Our results on the significance of *Unemployment*Forum* do not change (see Appendix A.46).

5.4. Partisan fear of full employment

Finally, we test our third hypothesis, that is that the fear of full employment varies based on which party controls the White House. We create a dummy variable, *President*, that takes the value of 1 when a Democrat is President and 0 when a Republican is President. Building on the models in Section 5.3, we study the effect that the Presidential party has when considering the differential reaction in FOMC meetings versus public speeches to changes in unemployment. In Table 3, we interact *President* with *Unemployment* and *Forum*. In doing this, we test the hypothesis that Fed officials react to changes in *Unemployment* based not only on *Forum*, but that these effects are also conditional on which party is in the White House.

The model results in Table 3 suggest that these Presidential effects are highly significant. The coefficient on *President* in model 1 is statistically significant and of considerable relative magnitude, and when testing interaction terms in models 2-4 we also see statistical significance on the coefficient for the two-way interaction between *Forum* and *President*, and the three-

[18]While it does not affect the significance of *Forum*, there is some interesting heterogeneity between Board members and regional presidents in Appendix A.45. While the interaction term *Unemp*Forum* is negative and significant for Board members (stronger reactions to low levels of unemployment in FOMC meetings), for regional presidents it is instead the main effect term *Forum* which is positive and significant (more *Labor* discussion across levels of unemployment in the FOMC). *Unemployment* remains highly significant in both.

way interaction including *Unemployment*. This suggests that *President* moderates the effect of *Forum* on how Fed officials react to changes in *Unemployment*. A joint Wald test confirms that presidential party is a significant determinant of Federal Reserve communication, with all six partisanship-related terms in model 4 jointly significant ($p < 0.001$). We again find no significance on the wage variables included in all regressions in Table 3. Given the challenges of interpreting the effects of this interaction directly from the table given the several interaction and quadratic terms, the results of the interaction model 4 are visualized in Figure 8 and Figure 9.

Figure 8 shows predictions from model 4 for the volume of discussion of *Labor* at different levels of *Unemployment*, comparing both across public speeches and FOMC meetings as well as when the president is a Democrat versus Republican. The predicted values are somewhat similar in public speeches across presidential party, with the shape of the curve slightly steeper when a Democrat is in the White House. However, the effect is much more pronounced in FOMC meetings, where at almost any level of unemployment the model predicts there will be more discussion of Labor as an inflation driver. Figure 9 plots the conditional effects of the president being a Democrat compared to a Republican in each forum, with the y-axis displaying the difference in Labor mentions due to this variable at a given level of unemployment. The lines are solid when the difference in Labor mentions based on *President* is statistically significant. While this difference is positive and significant at almost any level of unemployment in the FOMC, it is limited to significance only at low levels of unemployment in public speeches. This also provides evidence that there is not only a partisan fear of full employment across communication, but that this partisan fear is also genuine. It is in private policy deliberations, rather than public communications, where the partisan effect is strongest.

Table 3: Results on H3, a partisan fear of full employment, 1978-2019

	1	2	3	4
Unemployment	−0.143‡‡	−0.142‡‡	−0.047	−0.052
	(0.042)	(0.041)	(0.062)	(0.064)
Unemployment^2	0.056‡‡	0.055‡‡	0.057‡	0.062‡
	(0.017)	(0.017)	(0.021)	(0.021)
Forum (FOMC=1)	0.133	−0.120	−0.086	−0.106
	(0.107)	(0.125)	(0.140)	(0.139)
Unemp*Forum	−0.170‡‡	−0.194‡‡	−0.297‡‡	−0.311‡‡
	(0.049)	(0.048)	(0.064)	(0.063)
President (Dem=1)	0.411‡‡	0.159	0.164	0.216
	(0.077)	(0.101)	(0.135)	(0.134)
Pres*Forum		0.507‡‡	0.471†	0.453†
		(0.136)	(0.185)	(0.182)
Pres*Unemp			−0.190†	−0.166
			(0.087)	(0.089)
Unemp*Forum*Pres			0.234†	0.242†
			(0.099)	(0.098)
Num.Obs.	332	332	332	332
AIC	2316.2	2304.7	2305.9	2300.4
BIC	2380.9	2373.2	2378.2	2384.1
Chair dummies	✓	✓	✓	✓
Macro controls	✓	✓		✓
Log(total words)	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: Macro controls include real wage growth, inflation, and real GDP growth. As in other specifications, only GDP growth reaches significance. All interaction terms including *Unemployment^2* are included in model but do not have significant effects and are omitted from the table for simplicity.

Figure 8: Predictions, Model 4

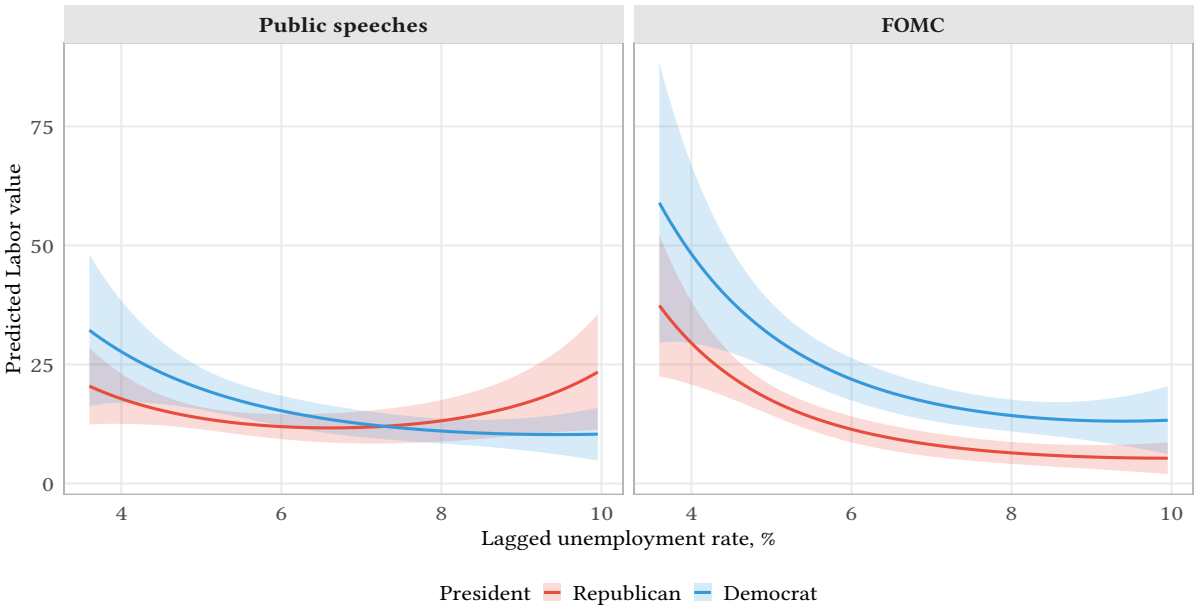

Note: Model predictions calculated while holding other variables at their means and averaging across chairs. Plot created using the marginaleffects package (Arel-Bundock et al., 2024).

Figure 9: Conditional effects, Model 4

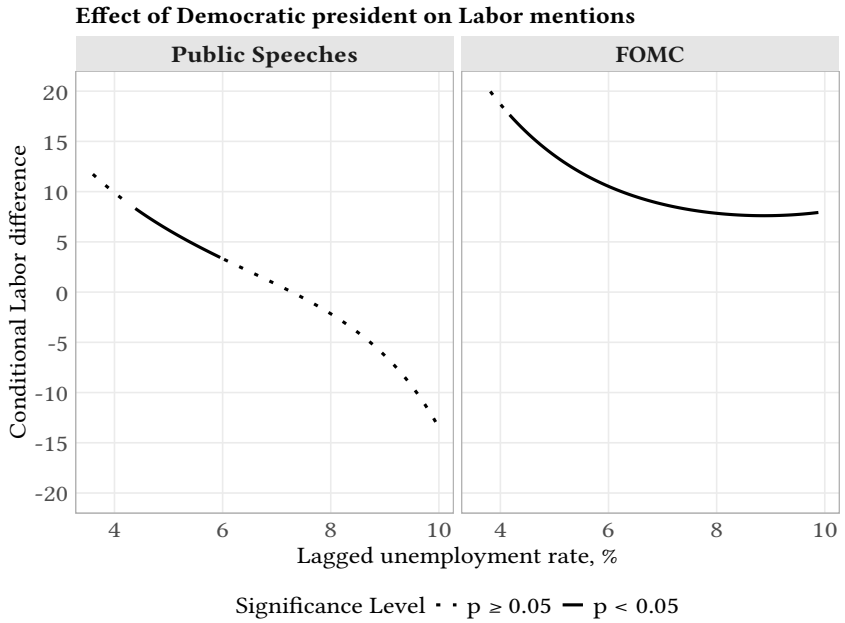

Note: Plot shows the conditional effect of the President being a Democrat (compared to a Republican) on Labor mentions, while holding other variables at their means and averaging across chairs. Plot created using the marginaleffects package (Arel-Bundock et al., 2024).

As in the other sections, we run a range of robustness checks on these regressions in Appendix A.5, including the use of Tealbook data and substituting the unemployment gap as

our *Unemployment* variable. We also perform regressions that include Congressional parties as well as presidential parties in Appendix A.53. While we find mixed results on the significance of Congressional party control, the effects of the president's party remain significant in all Congressional specifications, with consistent sign and comparable effect size. The effects of Congressional parties, however, are highly sensitive to the time period included in the regression.

5.5. Further considerations

Thus far in our analysis we have focused on the relationships between these variables for the entire period we have data available, 1978-2019, accommodating shifts in these relationships over time through the use of Chair dummies and other strategies like including an estimate of the unemployment gap. However, it is nonetheless possible that these relationships could have shifted over the time period of our sample in ways that are of empirical and theoretical interest. Particularly in the case of our analysis of partisan effects, it is also likely given the relatively small number of presidential administrations in our sample for each party that particular administrations may be influential in our results.

To probe our results, we conduct a leave-one-out analysis, iteratively re-estimating models 1 and 2 in Table 3 with one presidential administration excluded at a time. This sensitivity check enables us to examine to what extent our results are driven by a particular presidential administration. Figure 10 plots the coefficient estimates for the *Forum* and *Forum*Unemployment* variables from model 1 (which controls for the main effect from *President*), and Figure 11 plots the coefficient estimates for the *President* and *President*Forum* variables from model 2.[19] On the whole, our primary findings are largely consistent across iterations, with a couple of important observations.

First, our finding on the differential response to unemployment across FOMC meetings compared to public speeches is highly robust across each model in Figure 10 (*Unemployment*Forum*). The one exception is that when Reagan is dropped, the *Forum* main effect becomes significant at the 5% level, though the sign and direction of the *Unemployment*Forum* interaction remains consistent. These results are further corroborated by a leave-one-out analysis of Model 5 from Table 2 (without a presidential variable), dropping one Fed Chair at a time rather than president (see Figure 23). Together, these analyses suggest that the differential effect of *Forum* on Fed communication is not driven by any single period. The 1980s,

[19]Estimates for each of these variables are mostly consistent in sign and significance in a leave-one-out analysis of model 4. However, the three-way *Unemployment*Forum*President* interaction is less stable. While the sign is consistent in six out of seven runs, significance at the 5% level of the interaction term is lost when Reagan, Clinton, or Obama are removed.

however, warrant closer attention as an instructive exception. Given the dramatic variation in unemployment over that decade, it is unsurprising that the model's ability to detect the interaction weakens when Reagan-era observations are removed. Descriptive plots reveal a more nuanced picture: the period from approximately 1978 to 1982 is the only one in our sample where Labor discussion in public speeches consistently outstrips that in FOMC meetings, strikingly so during a period of high unemployment (see Figure 18). This likely reflects a deliberate communicative strategy by the Volcker Fed, which publicly framed wage and labor cost dynamics as the primary transmission mechanism of inflation, requiring sustained external justification as unemployment climbed (Mitchell & Erickson, 2005; Volcker, 1981).

Figure 10: Leave-one-out analysis, forum variables, Model 1

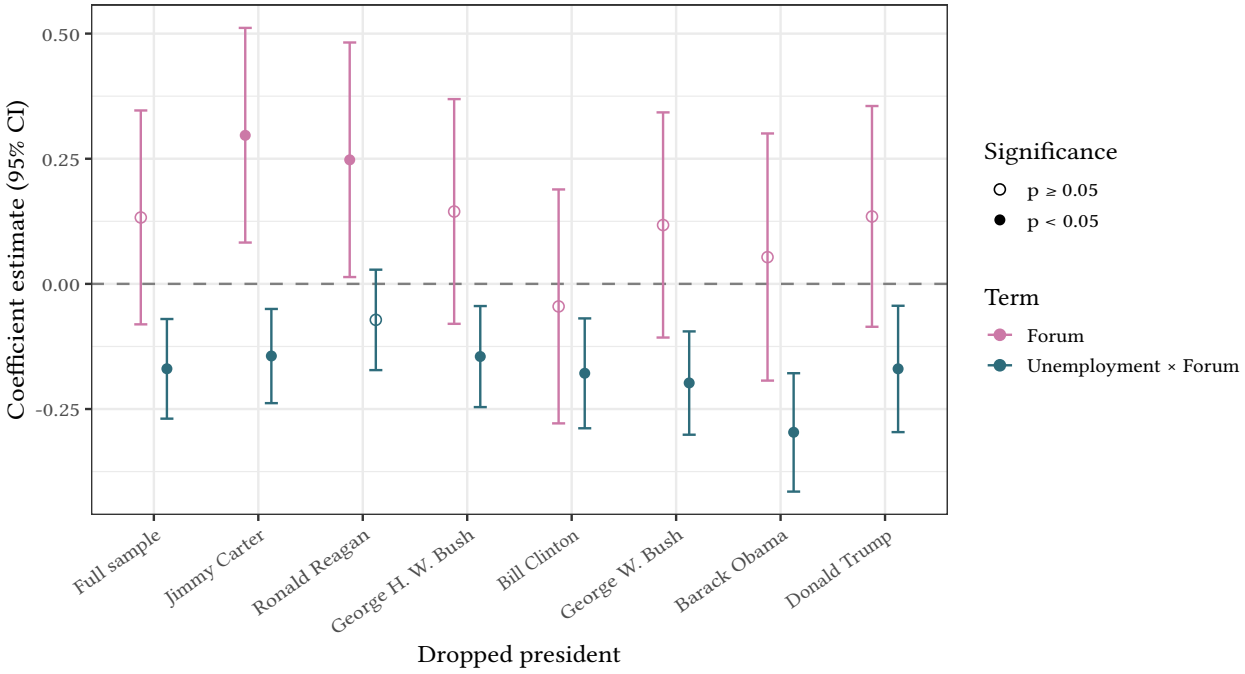

Second, the leave-one-out analysis also largely reinforces our finding that Fed officials' internal deliberations are more sensitive to partisan context than their public speech behavior. As in Figure 11, a significant partisan effect is maintained across our sample when dropping any given presidential administration. However, there is some meaningful heterogeneity when dropping certain administrations which is again worth reflecting on. There are two distinct dynamics at play: first, when Reagan or Obama are dropped, the *President* main effect becomes significant at the 5% level and the *President*Forum* term loses significance, though the sign remains consistent. Second, when Clinton is dropped, the interaction effect remains significant and positive, but the main effect of *President* switches sign, though does not achieve significance. These observations reflect the importance of each of these three sample periods

in our overall results. They demonstrate the influence of the Obama and Reagan presidencies in identifying the *President*Forum* effect, which is not unexpected given the long tenure of each of these administrations. Without either of these presidential administrations included, the partisan signal that was captured by the interaction term is instead reflected in the main effect. However, that the interaction estimate remains positive in all specifications even when insignificant demonstrates that the direction of the interaction is consistent, even if precision depends on the inclusion of these key administrations.

The effect of removing Clinton is especially informative to our findings, given that Clinton's presidency coincides with the most intensive period of *Labor* discussion in our sample in the late 1990s (as in Figure 18). That the *President*Forum* interaction remains positive and significant in the absence of the Clinton period is an important robustness check on how meaningful these results are across our entire sample period. The sign flip on the main effect for *President*, while not significant, reflects that the Clinton era's distinctive contribution to our overall findings is in the average volume of *Labor* communication under Democratic presidents, rather than in the forum-specific partisan contrast.

Figure 11: Leave one-out-analysis, presidential variables, Model 2

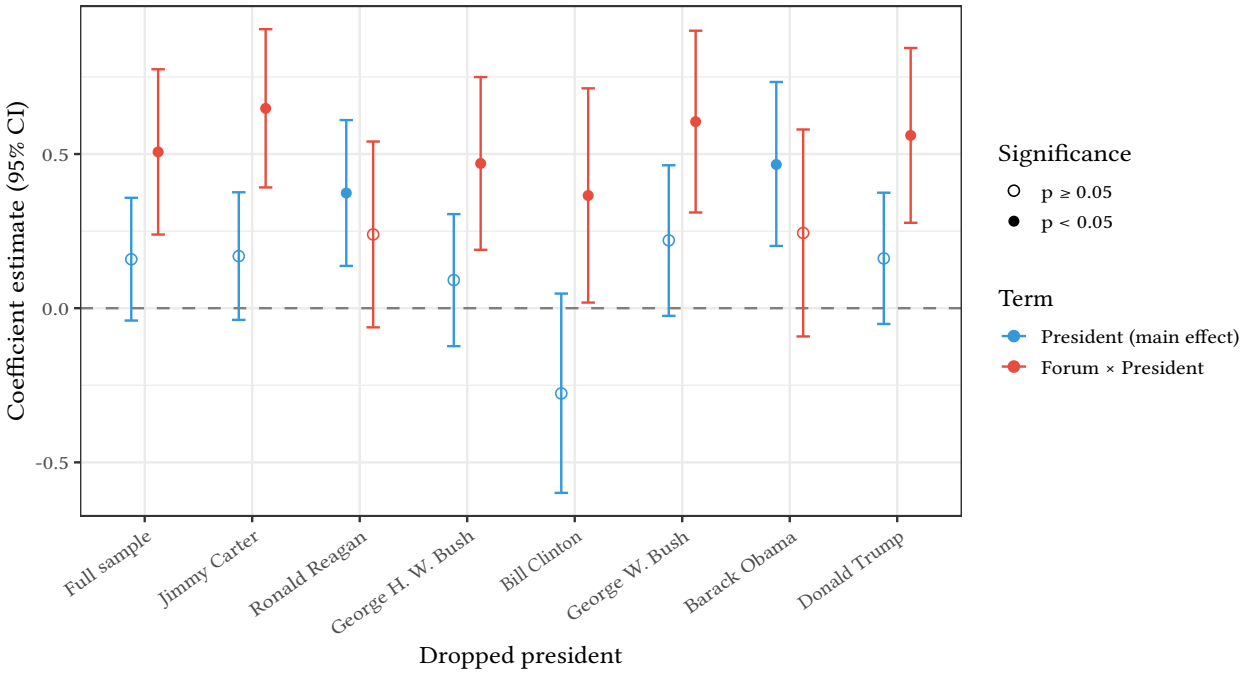

6. Conclusion

The Fed has a dual mandate for price stability and maximum employment. However, its *de facto* preference has been for the former ever since the Volcker Fed engineered a recession

and spike in unemployment to reign in inflation in the early 1980s. The theoretical basis for this preference was embodied in the expectation-augmented Phillips curve, positing a natural rate of unemployment below which inflation would rapidly spiral. While this arrangement has attracted myriad public, policymaker, and academic criticism, the Fed's ostensible "fear of full employment" has not been subjected to a systematic empirical test.

Our analysis fills this gap. Deploying LLMs to code Fed speeches and FOMC transcripts from 1978 to 2019, we construct a time series of instances where Fed officials point to labor market dynamics as drivers of inflation. We test three hypotheses of increasing stringency on the presence of a fear of full employment, with significant positive results on all three. First, we find that Fed officials place a greater emphasis on labor as a driver of inflation at lower levels of unemployment, regardless of actual growth in wages. Second, we document what we term *genuine* fear of full employment: that Fed officials are much more reactive to low levels of unemployment as inflationary in private policy deliberations (FOMC meetings) than in public speeches. Finally, we find evidence for a *partisan* fear of full employment: Fed officials place greater emphasis on labor as driving inflation when the president is a Democrat, holding constant macroeconomic conditions. These findings are robust to multiple model specifications, the use of alternative data sources, and a battery of other validation techniques.

We make several contributions to the literature. First, methodologically, we demonstrate how political scientists can apply LLMs in their work to generate new data which would have been previously too onerous to complete manually, and too complex for alternative natural language processing techniques. We employ a strict methodological pipeline which allows us to validate our results and use them as an input to our statistical analysis. Second, we offer new insights on the politics of full employment in the United States, finding strong empirical support for the charge that the Fed's monetary policy has been guided by a "baseless fear of unemployment" (Galbraith et al., 2007). We show that Fed policymakers have remained fearful of full employment despite decades of sluggish wage growth. This result bolsters the critique of the anti-worker bias of the inflation targeting regime in general, and of the Phillips curve and NAIRU concepts in particular (Arbogast et al., 2024; Galbraith, 1997). Given the importance of central bank choices for labor market outcomes, our findings also have broader relevance for a plethora of other effects linked to labor market experiences, from political attitudes to policy preferences and voting patterns (Bisbee & Rosendorff, 2025; Cox, 2024; Margalit, 2011; 2013; Pardos-Prado & Xena, 2019).

References

Acemoglu, D., & Restrepo, P. (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica*, *90*(5), 1973–2016. https://doi.org/10.3982/ECTA19815

Adolph, C. (2018). The Missing Politics of Central Banks. *PS: Political Science & Politics*, *51*(4), 737–742. https://doi.org/10.1017/S1049096518000847

Ainsley, C. (2021). Decentralized central banks: Political ideology and the Federal Reserve System of regional banks. *Governance*, *34*, 277–294.

Aklin, M., & Kern, A. (2021). The Side Effects of Central Bank Independence. *American Journal of Political Science*, *65*(4), 971–987. https://doi.org/10.1111/ajps.12580

Alvarez, M. J. A., Bluedorn, M. J. C., Hansen, M. N.-J. H., Huang, Y., Pugacheva, E., & Sollaci, A. (2022). Wage-Price Spirals: What is the Historical Evidence?. *IMF Working Papers*. https://ideas.repec.org//p/imf/imfwpa/2022-221.html

Arbogast, T., Van Doorslaer, H., & Vermeiren, M. (2024). Another strange non-death: the NAIRU and the ideational foundations of the Federal Reserve's new monetary policy framework. *Review of International Political Economy*, *31*(4), 805–830. https://doi.org/10.1080/09692290.2023.2250348

Arel-Bundock, V., Greifer, N., & Heiss, A. (2024). How to Interpret Statistical Models Using marginaleffects for R and Python. *Journal of Statistical Software*, *111*(9), 1–32.

Autor, D. H. (2015). Why Are There Still So Many Jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives*, *29*(3), 3–30. https://doi.org/10.1257/jep.29.3.3

Autor, D. H., Dorn, D., & Hanson, G. H. (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, *103*(6), 2121–2168. https://doi.org/10.1257/aer.103.6.2121

Bailey, S. K. (1950). *Congress Makes a Law: The Story Behind the Employment Act of 1946*. Columbia University Press.

Barrie, C., Palaiologou, E., & Törnberg, P. (2025). *Prompt Stability Scoring for Text Annotation with Large Language Models*.

Bartels, L. M. (2016). *Unequal democracy: The political economy of the new gilded age*.

Beland, L.-P. (2015). Political Parties and Labor-Market Outcomes: Evidence from US States. *American Economic Journal: Applied Economics*, *7*(4), 198–220. https://doi.org/10.1257/app.20120387

Bellodi, L. (2023). A Dynamic Measure of Bureaucratic Reputation: New Data for New Theory. *American Journal of Political Science*, *67*(4), 880–897. https://doi.org/10.1111/ajps.12695

Best, J. (2020). The quiet failures of early neoliberalism: From rational expectations to Keynesianism in reverse. *Review of International Studies*, *46*(5), 594–612. https://doi.org/10.1017/S0260210520000169

Beveridge, W. (1944). *Full employment in a free society*. George Allen & Unwin.

Binder, C. C. (2021). Political Pressure on Central Banks. *Journal of Money, Credit and Banking*, *53*(4), 715–744. https://doi.org/10.1111/jmcb.12772

Binder, S., & Spindel, M. (2017). *The myth of independence: How Congress governs the Federal Reserve*. Princeton University Press.

Bisbee, J., & Rosendorff, B. P. (2025). Antiglobalization sentiment: Exposure and immobility. *American Journal of Political Science*, *69*(3), 943–960.

Blinder, A. S. (2022). A Monetary and Fiscal History of the United States, 1961–2021. In *A Monetary and Fiscal History of the United States, 1961–2021: A Monetary and Fiscal History of the United States, 1961–2021*. Princeton University Press.

Blinder, A. S., & Watson, M. W. (2016). Presidents and the US Economy: An Econometric Exploration. *American Economic Review*, *106*(4), 1015–1045. https://doi.org/10.1257/aer.20140913

Blyth, M., & Fraccaroli, N. (2025). *Inflation: a guide for users and losers* (First edition). W. W. Norton & Company.

Braun, B. (2016). Speaking to the people? Money, trust, and central bank legitimacy in the age of quantitative easing. *Review of International Political Economy*, *23*(6), 1064–1092. https://doi.org/10.1080/09692290.2016.1252415

Campiglio, E., Deyris, J., Romelli, D., & Scalisi, G. (2025). Warning words in a warming world: Central bank communication and climate change. *European Economic Review*, *178*, 105101. https://doi.org/10.1016/j.euroecorev.2025.105101

Clark, W. R., & Arel-Bundock, V. (2012, October 24). *Independent but Not Indifferent: Partisan Bias in Monetary Policy at the Fed*. Social Science Research Network. https://papers.ssrn.com/abstract=2166647

Cova, J., & Schmitz, L. (2024). *A primer for the use of classifier and generative large language models in social science research* (OSF Preprints No. r3qng). https://econpapers.repec.org/paper/osfosfxxx/r3qng.htm

Cox, L. (2024). Great expectations: The effect of unmet labor market expectations after higher education on ideology. *American Journal of Political Science*, *68*(4), 1416–1430.

Cusack, T. R. (2001). Partisanship in the setting and coordination of fiscal and monetary policies. *European Journal of Political Research*, *40*(1), 93–115. https://doi.org/10.1111/1475-6765.00591

Diessner, S. (2023). The power of folk ideas in economic policy and the central bank–commercial bank analogy. *New Political Economy*, *28*(2), 315–328.

Do, S., Ollion, E., & Shen, R. (2022). The Augmented Social Scientist: Using Sequential Transfer Learning to Annotate Millions of Texts with Human-Level Accuracy. *Sociological Methods & Research*, 00491241221134526. https://doi.org/10.1177/00491241221134526

Downey, L. (2024). *Our Money: Monetary Policy as if Democracy Matters*. Princeton University Press.

Drechsler, I., Savov, A., & Schnabl, P. (2022). *Credit crunches and the great stagflation*.

Ferrara, F. M., Masciandaro, D., Moschella, M., & Romelli, D. (2021). Political voice on monetary policy: Evidence from the parliamentary hearings of the European Central Bank. *European Journal of Political Economy*, 102143. https://doi.org/10.1016/j.ejpoleco.2021.102143

Forder, J. (2010). Friedman's Nobel Lecture and the Phillips curve myth. *Journal of the History of Economic Thought*, *32*(3), 329–348.

Fraccaroli, N., Arel-Bundock, V., & Blyth, M. (2025). What do central bankers talk about when they talk about inflation? The rise and fall of inflation narratives. *New Political Economy*, *30*(5), 713–728. https://doi.org/10.1080/13563467.2025.2504392

Friedman, M. (1968). The Role of Monetary Policy. *The American Economic Review*, *58*(1), 1–17.

Galbraith, J. K. (1997). Time to Ditch the NAIRU. *Journal of Economic Perspectives*, *11*(1), 93–108.

Galbraith, J. K., Giovannoni, O., & Russo, A. J. (2007). *The Fed's real reaction function: monetary policy, inflation, unemployment, inequality and presidential politics* (No. UTIP Working Paper 42).

Garg, P., & Fetzer, T. (2025, January 14). *Causal Claims in Economics*. arXiv. https://doi.org/10.48550/arXiv.2501.06873

Goutsmedt, A. (2021). From stagflation to the Great Inflation: Explaining the US economy of the 1970s. *Revue d'economie politique*, *131*(3), 557–582. https://www.cairn-int.info/article.php?ID_ARTICLE=E_REDP_313_0239

Goutsmedt, A. (2022). How the Phillips Curve Shaped Full Employment Policy in the 1970s: The Debates on the Humphrey-Hawkins Act. *History of Political Economy*, *54*(4), 619–653.

Goutsmedt, A., & Fontan, C. (2024). The ECB and the inflation monsters: strategic framing and the responsibility imperative (1998–2023). *Journal of European Public Policy*, *31*(4), 999–1025. https://doi.org/10.1080/13501763.2023.2281583

Greenspan, A. (1999). *Remarks by Chairman Alan Greenspan: The American economy in a world context*. https://www.federalreserve.gov/boarddocs/speeches/1999/19990506.htm

Greider, W. (1987). *Secrets of the Temple: How the Federal Reserve Runs the Country*. Simon & Schuster.

Grimmer, J., Roberts, M. E., & Stewart, B. M. (2022). *Text as data: A new framework for machine learning and the social sciences*. Princeton University Press.

Hibbs, D. A. (1977). Political Parties and Macroeconomic Policy. *American Political Science Review*, *71*(4), 1467–1487. https://doi.org/10.2307/1961490

Jacobs, L., & King, D. (2016). *Fed power: How finance wins*. Oxford University Press.

Judge, B. (2023). Piercing the veil of monetarism: a decomposition of American inflation, 1970–1985. *New Political Economy*, *0*(0), 1–15. https://doi.org/10.1080/13563467.2023.2200243

Kalecki, M. (1943). Political Aspects of Full Employment. *The Political Quarterly*, *14*(4), 322–330. https://doi.org/10.1111/j.1467-923X.1943.tb01016.x

Kaya, A. (2022). The Federal Reserve's move to an explicit inflation target: incremental policy shifts in techno-political institutions. *Review of International Political Economy*, *29*(5), 1625–1649. https://doi.org/10.1080/09692290.2021.1934073

Kaya, A., Golub, S., Kuperberg, M., & Lin, F. (2019). The Federal Reserve's Dual Mandate and the Inflation-Unemployment Tradeoff. *Contemporary Economic Policy*, *37*(4), 641–651. https://doi.org/10.1111/coep.12422

Krippner, G. (2011). *Capitalizing on Crisis*. Harvard University Press.

Margalit, Y. (2011). Costly Jobs: Trade-related Layoffs, Government Compensation, and Voting in U.S. Elections. *American Political Science Review*, *105*(1), 166–188.

Margalit, Y. (2013). Explaining Social Policy Preferences: Evidence from the Great Recession. *American Political Science Review*, *107*(1), 80–103.

Mitchell, D. J., & Erickson, C. L. (2005). Not yet dead at the Fed: Unions, worker bargaining, and economy-wide wage determination. *Industrial Relations: A Journal of Economy and Society*, *44*(4), 565–606.

Moschella, M. (2024). *Unexpected revolutionaries: how central banks made and unmade economic orthodoxy*. Cornell University Press.

Pagliuca, A. (2025, June 25). *Partisan Fed*. Social Science Research Network. https://doi.org/10.2139/ssrn.5320940

Pardos-Prado, S., & Xena, C. (2019). Skill Specificity and Attitudes toward Immigration. *American Journal of Political Science*, *63*(2), 286–304.

Phelps, E. S. (1967). Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time. *Economica*, *34*(135), 254–281.

Phelps, E. S. (1968). Money-Wage Dynamics and Labor-Market Equilibrium. *The Journal of Political Economy*, *76*(4), 678–711.

Phillips, A. (1958). The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957. *Economica*, *25*(100), 283–299.

Piketty, T., & Saez, E. (2003). Income Inequality in the United States, 1913-1998. *The Quarterly Economic Journal*, *118*(1), 1–39.

Powell, J. H. (2020,). Opening remarks: New economic challenges and the fed's monetary policy review. *Economic Policy Symposium, Jackson Hole*.

Roberts, J. M. (1995). New Keynesian Economics and the Phillips Curve. *Journal of Money, Credit and Banking*, *27*(4), 975–984. https://doi.org/10.2307/2077783

Romelli, D. (2022). The political economy of reforms in Central Bank design: evidence from a new dataset. *Economic Policy*, *37*(112), 641–688. https://doi.org/10.1093/epolic/eiac011

Romelli, D. (2024). Trends in central bank independence: a de-jure perspective. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4716704

Samuelson, P. A., & Solow, R. M. (1960). Analytical Aspects of Anti-Inflation Policy. *The American Economic Review*, *50*(2), 177–194. http://www.jstor.org/stable/1815021

Törnberg, P. (2024, February 5). *Best Practices for Text Annotation with Large Language Models*. arXiv. https://doi.org/10.48550/arXiv.2402.05129

Volcker, P. A. (1981). *Statement Before the Committee on Banking, Housing, and Urban Affairs*.

Volcker, P. A., & Harper, C. (2018). *Keeping at it: the quest for sound money and good government*. Hachette.

Weir, M. (1987). Full Employment as a Political Issue in the United States. *Social Research*, *54*(2), 377–402. https://www.jstor.org/stable/40970462

A Appendix

A.1 Additional information on large language models

A.11 Keyword list of inflation related terms

The 16 expressions used to identify inflation-related sentences are:

"inflation", "cpi", "pce", "price increas", "price stab", "overall price", "price level", "consumer price", "producer price", "price index", "price pressure", "price shock", "higher price", "raise price", "raising price", "rising price".

We first identified basic inflation expressions (inflation, cpi, pce), then examined the top 50 bigrams containing "price" in our corpus and retained all salient terms, to minimize false positives while being certain to not miss any relevant discussions.

A.12 Full codebook

- "*Labor*" if the excerpt states that wages, workers' demands, or other developments in labor markets cause inflation.
- "*Market power*" if the excerpt states that business' price-setting behaviour, profits, markups or competition dynamics cause inflation.
- "*Fiscal*" if the excerpt states that government spending, deficit, public expenses, taxes, and other fiscal or budgetary considerations cause inflation.
- "*Energy*" if the excerpt states that changes in the price or supply of energy cause inflation.
- "*Commodities*" if the excerpt states that changes in the price of non-energy commodities (such as food, agriculture, raw materials, etc.) cause inflation.
- "*Exchange rate*" if the excerpt states that the value of the dollar against other currencies causes inflation.
- "*Other*" if the excerpt discusses other causes of inflation that are not captured by the previous categories (such as supply chain bottlenecks, regulations, productivity, credit dynamics, shelter, healthcare costs, or anything not listed here).

Excerpts can have zero labels (when inflation is mentioned without discussing its drivers), or multiple labels (when several causes are discussed in turn).

A.13 Detailed model performance metrics: o4-mini, Llama 3.1 405B, and DeepSeek-V3

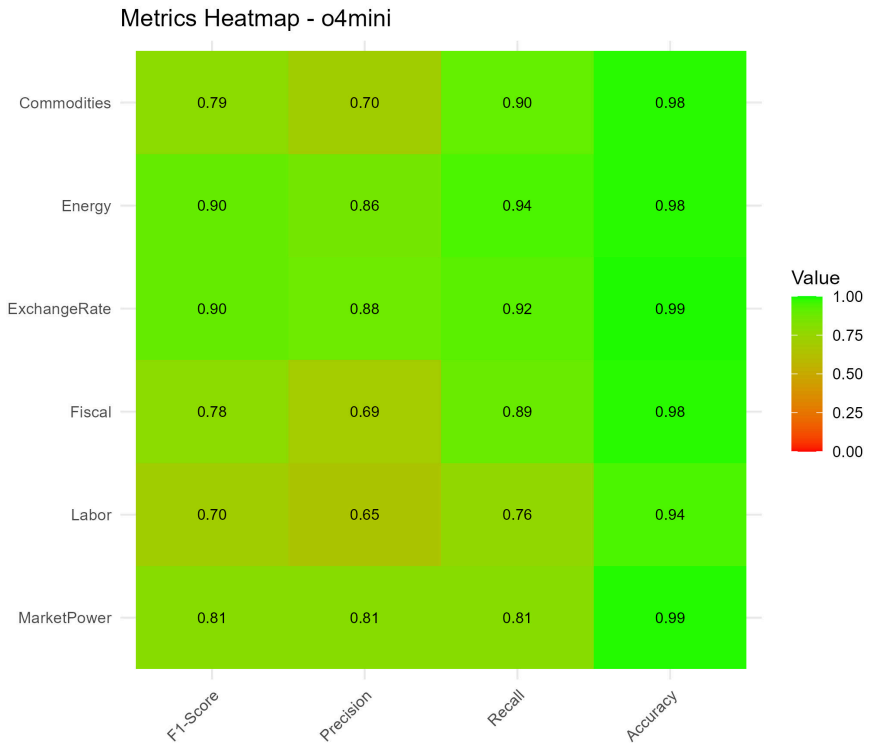

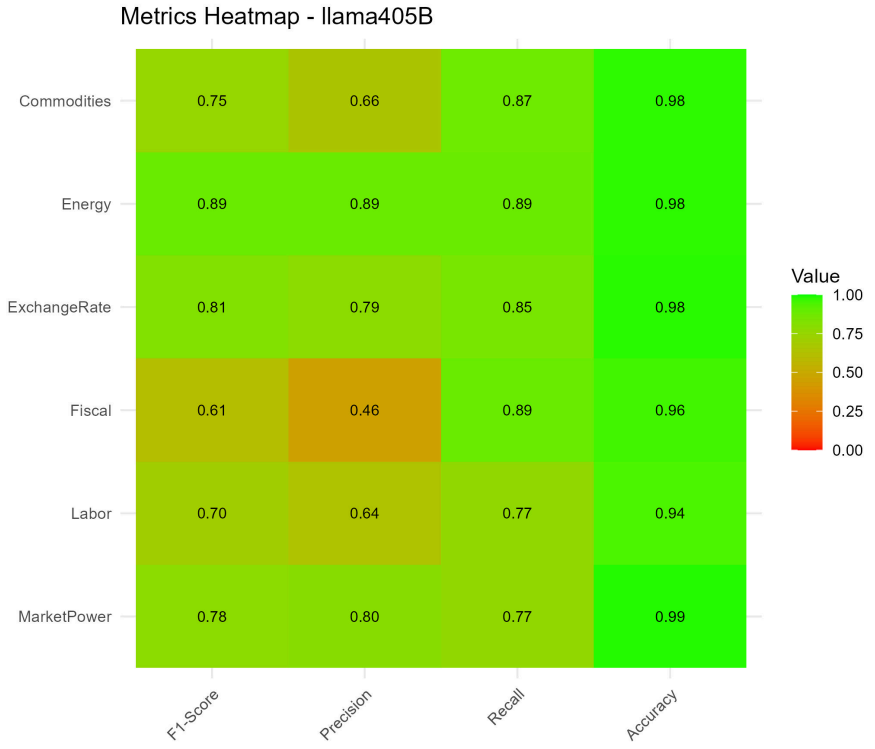

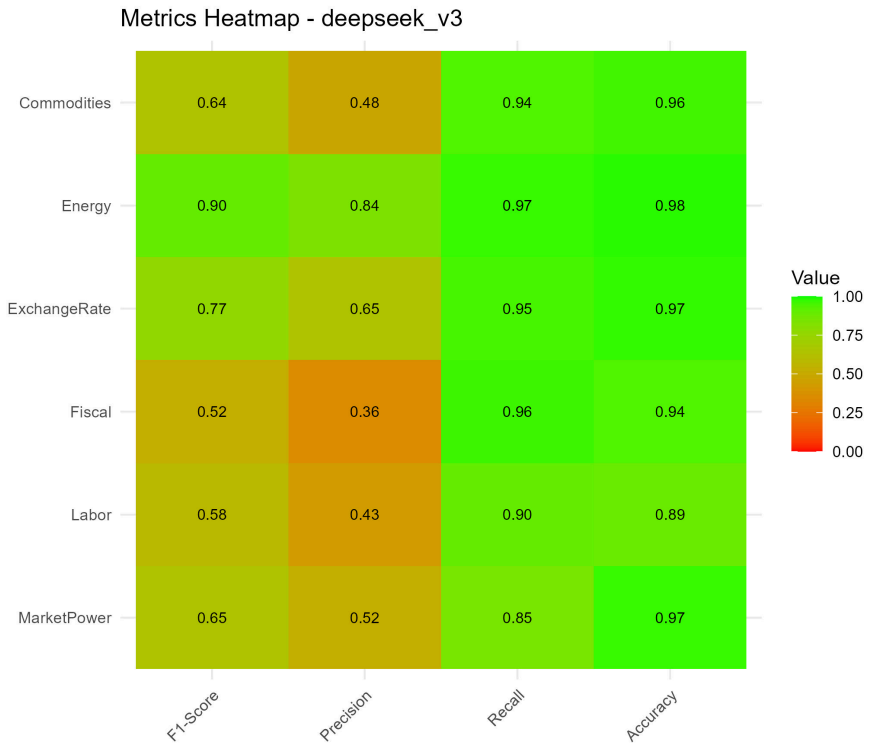

A.14 Prompt stability results

Figure 15: o4 mini

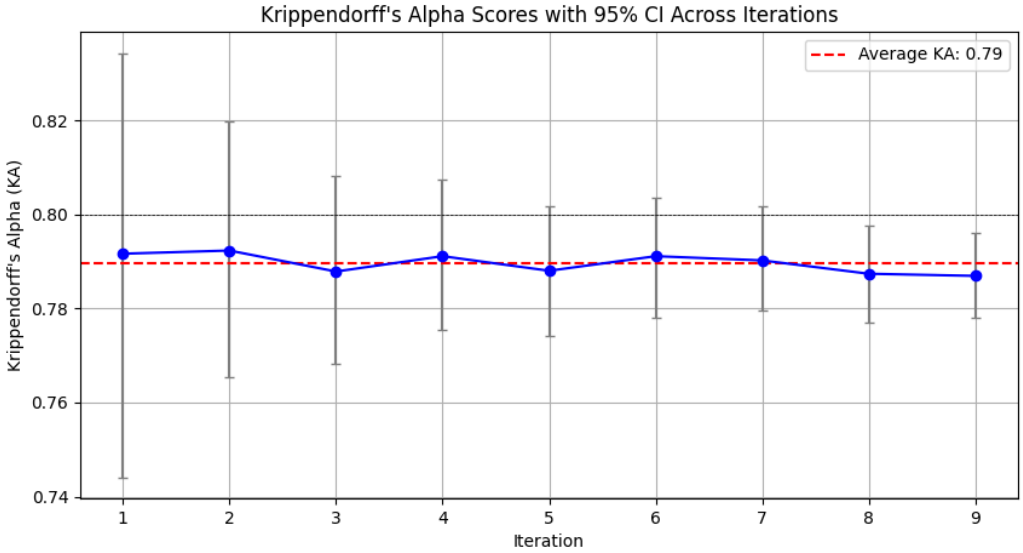

Figure 16: o3 mini

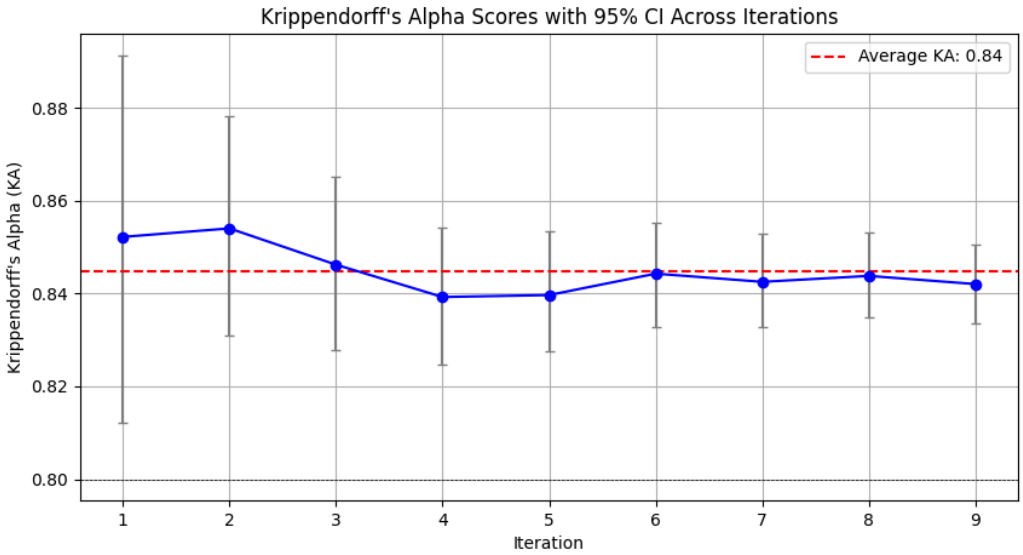

A.15 Final prompt to LLMs

```
We are exploring central bank communication by looking at central banker
policy deliberations. We want to understand how they discuss the causes of
inflation.
    Please read the excerpt, and classify it using the following labels:

- "None" if the excerpt does not discuss inflation.
- “None” if the excerpt describes inflation but does not make any reference
to the causes of inflation.
- “Labor” if the excerpt states that wages, workers' demands, or other
developments in labor markets cause inflation.
- “Market power” if the excerpt states that business' price-setting
behaviour, profits, markups or competition dynamics cause inflation.
- “Fiscal” if the excerpt states that government spending, deficit, public
expenses, taxes, and other fiscal or budgetary considerations cause inflation.
- “Energy” if the excerpt states that changes in the price or supply of
energy cause inflation.
- “Commodities" if the excerpt states that changes in the price of non-energy
commodities (such as food, agriculture, raw materials, etc.) cause inflation.
- “Exchange rate" if the excerpt states that the value of the dollar against
other currencies causes inflation.
- “Other" if the excerpt discusses other causes of inflation that are not
captured by the previous categories (such as supply chain bottlenecks,
regulations, productivity, credit dynamics, shelter, healthcare costs, or
anything not listed here).
```

```
You can give excerpts multiple labels (for example, “Market power, Fiscal”).

IF the excerpt has the label "None", it should have no other labels.

Reply ONLY with the assigned label(s) (for example, "Labor, Other"). DO NOT
EXPLAIN YOUR ANSWER. This is the excerpt:
```

A.2 Details on data sources for variables

Tealbook (formerly Greenbook) variables

Tealbooks are produced by staff at the Board of Governors, and are provided to FOMC members in advance of FOMC meetings, often one to two weeks ahead. Prior to 2010, detailed analysis of economic conditions was instead contained in Greenbooks (which merged with Bluebooks in June 2010 to become the Tealbook). Like FOMC transcripts, they are released publicly with a five-year lag. The Philadelphia Fed has consolidated the forecast values from historical Tealbooks and Greenbooks into a dataset that can be download from their website, which we use as our source (link). We refer to this data from Tealbooks and Greenbooks collectively here (and throughout the paper) as “Tealbooks”.

All Tealbooks contain “nowcasts” (forecasts) for the current quarter. They also contain historical values (which, once available, are the actual values drawn from the BEA or BLS), as well as forecasts for future quarters. To obtain our quarterly variables for our economic indicators, we average these “nowcast” values for a given quarter. For our unemployment variable, we use “UNEMP” values from the Philadelphia Fed dataset (unemployment rate). For our real GDP growth variable, we use “gRGDP” (quarter-over-quarter growth in real GDP, annualized percentage points). For our inflation variable, we use “gPGDP” (quarter-over-quarter growth in price index for GDP, annualized percentage points). Alternative Greenbook inflation metrics (CPI and PCE) are not available for our entire sample period.

Historical macroeconomic variables

In addition to the Greenbook “nowcasts”, we also use actual historical values for macroeconomic variables. We download the dataset for the unemployment rate from the U.S. Bureau of Labor Statistics (BLS) via FRED. Unemployment data from the BLS is typically released for each month on the first Friday of the month following. We average monthly unemployment figures to create a quarterly rate.

We download data for the percentage change in the Personal Consumption Expenditures (PCE) index (quarter-over-quarter change, annualized percentage points) and real GDP growth (quarter-over-quarter change, annualized percentage points) from the U.S. Bureau of

Economic Analysis (BEA) via DBnomics. For both, the first estimate for the quarter is usually released by the BEA at the end of the first month of the following quarter.[20]

Wage variables

We use two different metrics of wage growth in our analysis, Average Hourly Earnings (AHE, "Average Hourly Earnings of Production and Nonsupervisory Employees, Total Private") and Median Weekly Earnings (MWE, "Employed full time: Median usual weekly real earnings: Wage and salary workers: 16 years and over"). MWE is only available beginning in 1979. Both are released by the BLS, and downloaded via FRED. MWE is released quarterly in the month following the end of the quarter as part of the Current Population Survey. AHE is released monthly, usually around the mid-point of the month following, as part of the Current Employment Statistics. MWE is available for download on FRED in both nominal and real (1982-84 CPI Adjusted Dollars) units. AHE is only available to download in nominal units; we deflate nominal wages using the CPI (CPIAUCSL on FRED) to convert values into constant 1982-1984 dollars for our real AHE variable. To convert AHE from monthly to quarterly, we average values across the quarter. We then take quarter-over-quarter growth rates, annualized percentage points, for both AHE and MWE.

[20]There are also monthly releases of PCE data, in addition to a quarterly figure being included in the GDP release. These monthly releases also usually happen at the end of the month following.

A.3 Robustness checks for models in Section 5.2

A.31 Full results for specifications in main text

Table 4: Results on H1, a fear of full employment, 1978-2019

	1	2	3	4	5
Unemployment	−0.202‡‡	−0.202‡‡	−0.186‡‡	−0.196‡‡	−0.187‡‡
	(0.038)	(0.038)	(0.041)	(0.039)	(0.039)
Unemployment^2	0.060‡‡	0.061‡‡	0.059‡‡	0.059‡‡	0.056‡‡
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Real wage growth				0.015	0.020
				(0.033)	(0.032)
Nominal wage growth			0.033		
			(0.034)		
Inflation		0.007		0.020	0.033
		(0.021)		(0.035)	(0.035)
Real GDP growth					0.036†
					(0.014)
Quarterly word count (log)	0.506‡	0.511‡	0.510‡	0.510‡	0.516‡
	(0.172)	(0.172)	(0.171)	(0.173)	(0.169)
Num.Obs.	166	166	166	166	166
AIC	1346.3	1348.2	1347.4	1350.0	1345.7
BIC	1377.4	1382.4	1381.6	1387.4	1386.2
Chair dummies	✓	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

A.32 Alternative modelling strategies

To ensure that the negative binomial specification is the optimal choice for modelling our data, we also run model 5 from Section 5.2 using three other types of models commonly used with count data: Poisson, Poisson with sandwich standards errors, and quasi-Poisson.

Table 5: Results on H1, a fear of full employment, 1978-2019

	NB	Poisson	Poisson Sandwich	Quasi Poisson
Unemployment	−0.187‡‡	−0.172‡‡	−0.172‡‡	−0.172‡‡
	(0.039)	(0.014)	(0.033)	(0.038)
Unemployment^2	0.056‡‡	0.046‡‡	0.046‡	0.046‡
	(0.015)	(0.006)	(0.016)	(0.016)
Real wage growth	0.020	0.023	0.023	0.023
	(0.032)	(0.013)	(0.033)	(0.037)
Inflation	0.033	0.016	0.016	0.016
	(0.035)	(0.015)	(0.033)	(0.043)
Real GDP growth	0.036†	0.050‡‡	0.050‡	0.050‡
	(0.014)	(0.006)	(0.019)	(0.016)
Quarterly word count (log)	0.516‡	0.486‡‡	0.486‡	0.486‡
	(0.169)	(0.063)	(0.153)	(0.176)
Num.Obs.	166	166	166	166
AIC	1345.7	1982.7	1982.7	
BIC	1386.2	2020.0	2020.0	
Chair dummies	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: Model 5 from main text; "NB" is negative binomial.

As can be seen in Table 5, our results from the main text (column 1) hold with other model types. However, the negative binomial specification offers a superior fit when comparing AIC and BIC values, as well as the −2 x log likelihood figures in Table 6.

Table 6: −2x log likelihoods

Model	2 x Log-Likelihood
NB	−1319.74
Poisson	−1958.66

Note: "NB" is negative binomial.

A.33 Tealbook data

We also test the specification in Section 5.2 using Tealbook (formerly Greenbook) forecasts for a given quarter. Unlike the models using actual data sources, we do not lag Tealbook variables, given that they are intended to give Fed policymakers a contemporaneous idea of what macroeconomic indicators are likely to be for a given quarter ahead of policy meetings. More detail on these variables is available in Appendix A.2. The results in Table 7 are consistent with those in Table 4.

Table 7: Results on H1, a fear of full employment, 1978-2019

	1	2	3	4	5
Unemployment (GB)	−0.240‡‡	−0.239‡‡	−0.231‡‡	−0.242‡‡	−0.227‡‡
	(0.037)	(0.037)	(0.040)	(0.037)	(0.036)
Unemployment^2 (GB)	0.068‡‡	0.070‡‡	0.067‡‡	0.071‡‡	0.065‡‡
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Real wage growth (AHE)				−0.004	−0.001
				(0.021)	(0.020)
Nom wage growth (AHE)			0.021		
			(0.032)		
Inflation (GB)		0.019		0.017	0.065†
		(0.026)		(0.028)	(0.029)
GDP growth (GB)					0.080‡‡
					(0.022)
Quarterly word count (log)	0.529‡	0.528‡	0.497‡	0.496‡	0.478‡
	(0.165)	(0.165)	(0.165)	(0.166)	(0.159)
Num.Obs.	167	167	166	166	166
AIC	1344.6	1346.1	1335.5	1337.4	1326.8
BIC	1375.8	1380.4	1369.8	1374.8	1367.3
Chair dummies	✓	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: "GB" stands for Greenbook. Greenbook variables are not lagged; wage variables are lagged by one quarter to account for release dates.

A.34 Alternative wage data

Table 8: Results on H1, a fear of full employment, 1979-2019, alt. wage metric

	1	2	3	4	5
Unemployment	−0.202‡‡	−0.202‡‡	−0.206‡‡	−0.206‡‡	−0.197‡‡
	(0.038)	(0.038)	(0.039)	(0.039)	(0.038)
Unemployment^2	0.060‡‡	0.061‡‡	0.061‡‡	0.062‡‡	0.059‡‡
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Real wage growth (MWE)				−0.009	−0.005
				(0.013)	(0.013)
Nominal wage growth (MWE)			−0.007		
			(0.013)		
Inflation		0.007		0.001	0.012
		(0.021)		(0.024)	(0.024)
Real GDP growth					0.037†
					(0.015)
Quarterly word count (log)	0.506‡	0.511‡	0.511‡	0.519‡	0.520‡
	(0.172)	(0.172)	(0.173)	(0.173)	(0.169)
Num.Obs.	166	166	162	162	162
AIC	1346.3	1348.2	1315.0	1316.7	1312.7
BIC	1377.4	1382.4	1345.9	1350.7	1349.7
Chair dummies	✓	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: “MWE” stands for Median Weekly Earnings. MWE data is only available from 1979-Q2; with lag, analysis begins in 1979-Q3.

A.35 UGAP regressions

Figure 17: Comparison of UGAP and unemployment variables

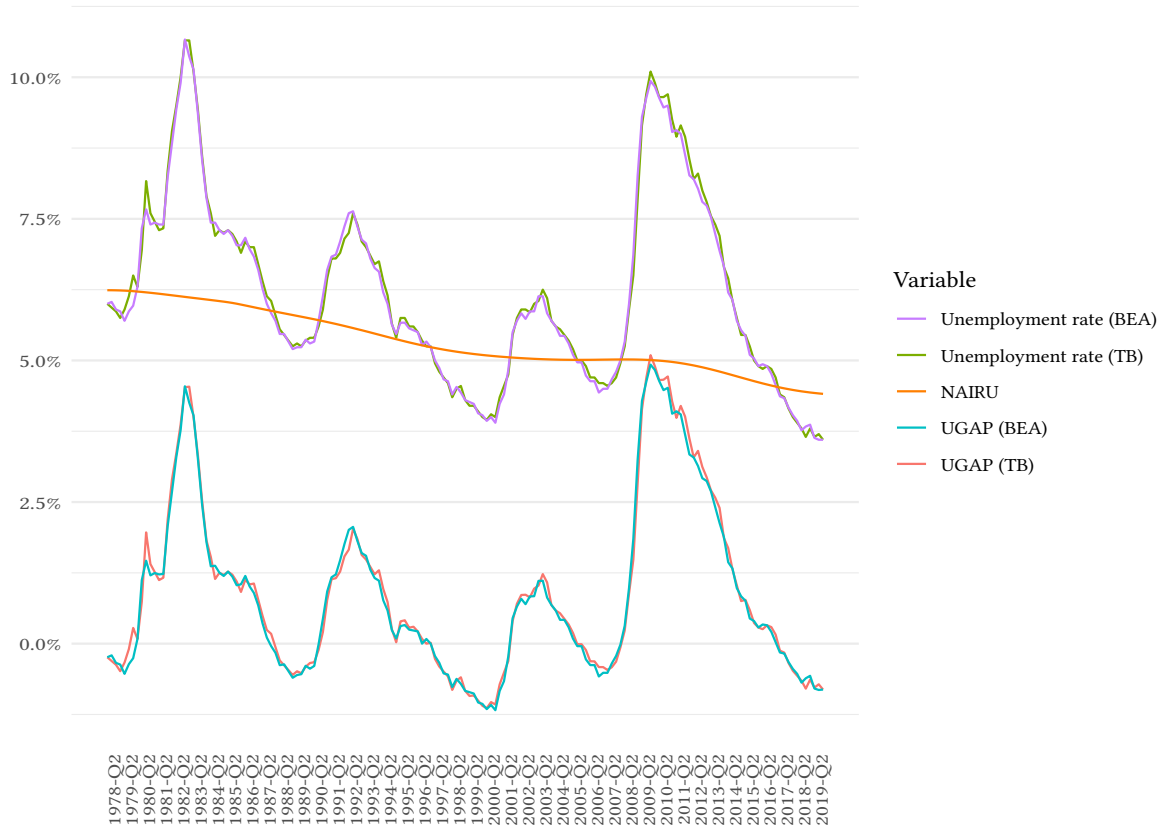

Note: "TB" stands for Tealbook. NAIRU rates drawn from Congressional Budget Office estimates.

Table 9: Results on H1, a fear of full employment, 1978-2019

	1	2	3	4	5
UGAP	−0.314‡‡	−0.317‡‡	−0.291‡‡	−0.296‡‡	−0.280‡‡
	(0.068)	(0.070)	(0.074)	(0.071)	(0.070)
UGAP^2	0.062‡‡	0.062‡‡	0.059‡	0.058‡	0.054‡
	(0.018)	(0.018)	(0.019)	(0.019)	(0.018)
Real wage growth				0.032	0.037
				(0.033)	(0.033)
Nominal wage growth			0.031		
			(0.035)		
Inflation		−0.009		0.019	0.034
		(0.022)		(0.037)	(0.036)
Real GDP growth					0.037†
					(0.015)
Quarterly word count (log)	0.549‡	0.544‡	0.551‡	0.540‡	0.545‡
	(0.175)	(0.176)	(0.175)	(0.176)	(0.172)
Num.Obs.	166	166	166	166	166
AIC	1354.0	1355.8	1355.3	1357.0	1352.7
BIC	1385.1	1390.1	1389.5	1394.4	1393.2
Chair dummies	✓	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: UGAP estimates using BEA unemployment data and NAIRU rates drawn from Congressional Budget Office estimates lagged by one quarter.

A.4 Robustness checks for models in Section 5.3

A.41 Time series by forum

Figure 18: Labor salience, 1978-2019

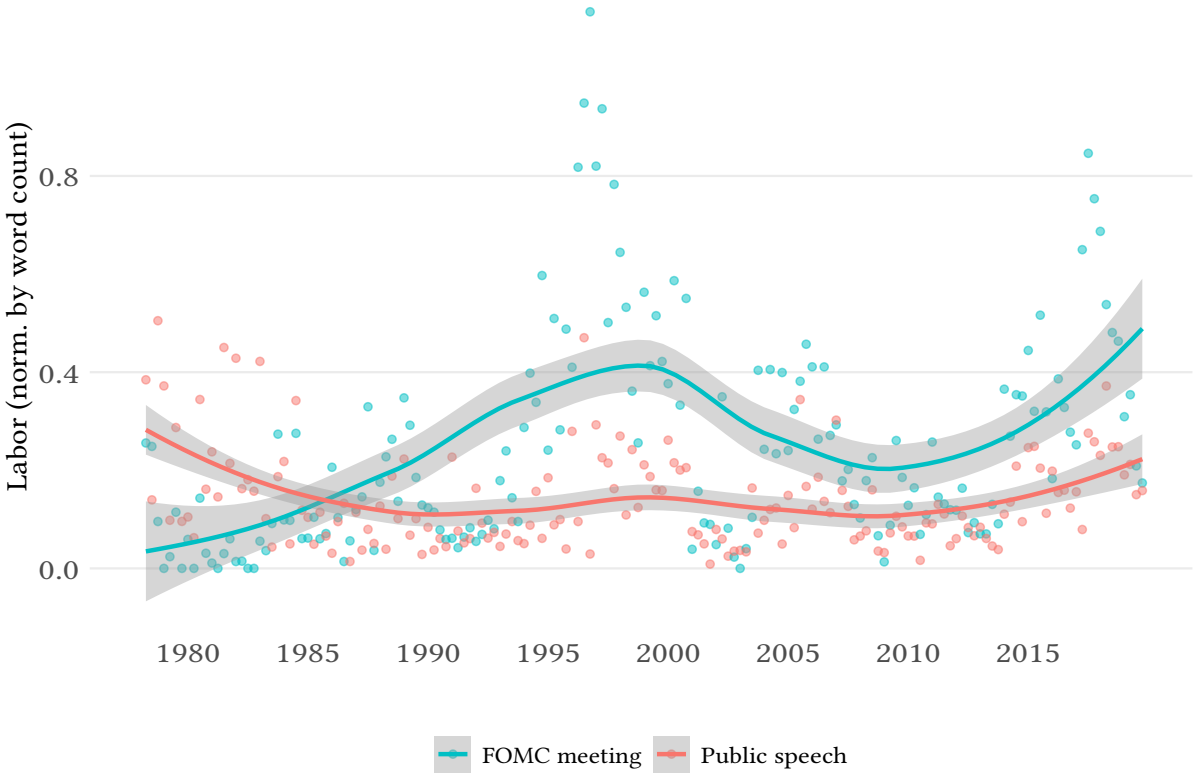

A.42 Tealbook data

Table 10: Results on H2, a genuine fear of full employment, 1978-2019

	1	2	3	4	5
Unemployment (GB)	−0.158‡‡	−0.158‡‡	−0.153‡‡	−0.163‡‡	−0.142‡‡
	(0.041)	(0.041)	(0.043)	(0.041)	(0.040)
Unemployment^2 (GB)	0.070‡‡	0.071‡‡	0.070‡‡	0.073‡‡	0.067‡‡
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
Forum (FOMC=1)	0.208	0.215†	0.217†	0.218†	0.202
	(0.109)	(0.109)	(0.110)	(0.110)	(0.107)
Unemp*Forum	−0.173‡‡	−0.171‡‡	−0.168‡‡	−0.169‡‡	−0.179‡‡
	(0.048)	(0.048)	(0.048)	(0.048)	(0.047)
Unemp^2*Forum	−0.006	−0.007	−0.009	−0.009	−0.010
	(0.022)	(0.022)	(0.022)	(0.022)	(0.021)
Real wage growth (AHE)				−0.006	−0.003
				(0.019)	(0.018)
Nominal wage growth (AHE)			0.019		
			(0.029)		
Inflation (GB)		0.017		0.015	0.064†
		(0.024)		(0.025)	(0.027)
Real GDP growth (GB)					0.086‡‡
					(0.020)
Quarterly word count (log)	0.478‡‡	0.479‡‡	0.467‡‡	0.467‡‡	0.456‡‡
	(0.098)	(0.098)	(0.098)	(0.098)	(0.095)
Num.Obs.	334	334	332	332	332
AIC	2341.4	2343.0	2324.2	2326.0	2310.5
BIC	2391.0	2396.3	2377.5	2383.1	2371.4
Chair dummies	✓	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: "GB" stands for Greenbook. Greenbook variables are not lagged; wage variables are lagged by one quarter to account for release dates.

A.43 Alternative wage data

Table 11: Results on H1, a fear of full employment, 1979-2019, alt. wage metric

	1	2	3	4	5
Unemployment	−0.128‡	−0.128‡	−0.138‡	−0.138‡	−0.126‡
	(0.042)	(0.043)	(0.042)	(0.042)	(0.041)
Unemployment^2	0.061‡‡	0.061‡‡	0.065‡‡	0.066‡‡	0.064‡‡
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
Forum (FOMC=1)	0.194	0.197	0.250†	0.254†	0.233†
	(0.112)	(0.112)	(0.112)	(0.112)	(0.111)
Unemp*Forum	−0.158‡	−0.158‡	−0.149‡	−0.149‡	−0.155‡
	(0.051)	(0.051)	(0.050)	(0.050)	(0.050)
Unemp^2*Forum	−0.002	−0.003	−0.011	−0.012	−0.012
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
Real wage growth (MWE)				−0.009	−0.006
				(0.012)	(0.012)
Nominal wage growth (MWE)			−0.007		
			(0.011)		
Inflation		0.006		−0.000	0.010
		(0.019)		(0.022)	(0.021)
Real GDP growth					0.037‡
					(0.014)
Quarterly word count (log)	0.452‡‡	0.454‡‡	0.452‡‡	0.455‡‡	0.437‡‡
	(0.100)	(0.101)	(0.099)	(0.099)	(0.098)
Num.Obs.	332	332	324	324	324
AIC	2342.7	2344.6	2278.3	2279.9	2274.3
BIC	2392.2	2397.9	2327.4	2332.8	2331.0
Chair dummies	✓	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: "MWE" stands for Median Weekly Earnings. MWE data is only available from 1979-Q2; with lag, analysis begins in 1979-Q3.

A.44 Unemployment gap

Table 12: Results on H2, a genuine fear of full employment, 1978-2019

	1	2	3	4	5
UGAP	−0.254‡	−0.257‡	−0.234‡	−0.241‡	−0.221‡
	(0.082)	(0.083)	(0.086)	(0.085)	(0.084)
UGAP^2	0.058‡	0.058‡	0.057†	0.055†	0.051†
	(0.022)	(0.022)	(0.022)	(0.022)	(0.022)
Forum (FOMC=1)	0.274‡	0.271†	0.280‡	0.273‡	0.251†
	(0.105)	(0.105)	(0.105)	(0.105)	(0.104)
UGAP*Forum	−0.139	−0.139	−0.131	−0.134	−0.143
	(0.111)	(0.111)	(0.111)	(0.110)	(0.109)
UGAP^2*Forum	0.010	0.010	0.008	0.010	0.011
	(0.030)	(0.030)	(0.030)	(0.030)	(0.029)
Real wage growth				0.029	0.033
				(0.030)	(0.030)
Nominal wage growth			0.031		
			(0.032)		
Inflation		−0.009		0.016	0.028
		(0.020)		(0.033)	(0.033)
Real GDP growth					0.035‡
					(0.013)
Quarterly word count (log)	0.436‡‡	0.433‡‡	0.439‡‡	0.432‡‡	0.413‡‡
	(0.103)	(0.103)	(0.103)	(0.103)	(0.102)
Num.Obs.	332	332	332	332	332
AIC	2361.2	2363.0	2362.3	2364.3	2359.7
BIC	2410.7	2416.3	2415.6	2421.3	2420.5
Chair dummies	✓	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: UGAP estimates using BEA unemployment data and NAIRU rates drawn from Congressional Budget Office estimates lagged by one quarter.

A.45 Speaker distribution

Figure 19: Total Fed Corpus by Position, Percentage

Note: Includes entire corpus (all inflation and non-inflation speech)

Figure 20: Labor salience by forum and speaker type

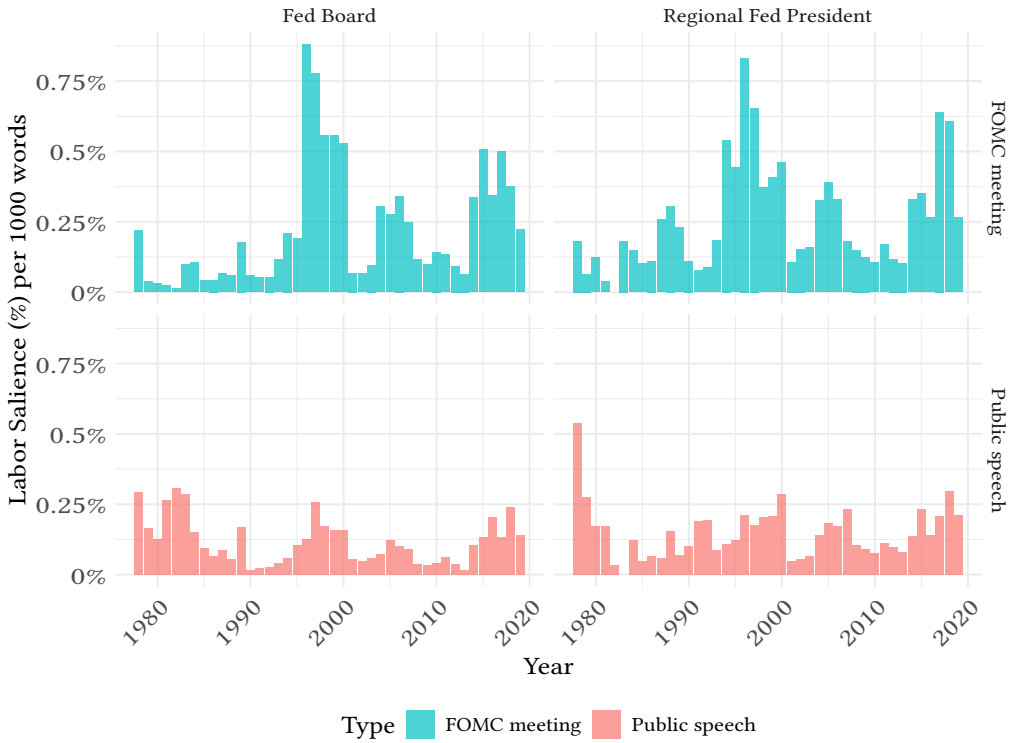

Table 13: Results for H2, subgroup regression

	Fed Board		Regional Presidents	
	1	2	3	4
Unemployment	−0.171‡	−0.131†	−0.142‡‡	−0.156‡‡
	(0.059)	(0.059)	(0.042)	(0.043)
Unemployment^2	0.088‡‡	0.072‡	0.013	0.013
	(0.023)	(0.024)	(0.020)	(0.020)
Forum (FOMC=1)	0.236	0.218	0.336‡	0.307‡
	(0.152)	(0.150)	(0.110)	(0.108)
Unemp*Forum	−0.219‡	−0.218‡	−0.060	−0.065
	(0.071)	(0.070)	(0.051)	(0.050)
Unemp^2*Forum	0.012	0.013	0.014	0.018
	(0.031)	(0.031)	(0.025)	(0.024)
Num.Obs.	332	332	332	332
AIC	1842.3	1840.3	1978.2	1968.7
BIC	1891.8	1901.2	2027.7	2029.6
Chair dummies	✓	✓	✓	✓
Macro controls		✓		✓
Log(total words)	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: Fed Board includes all members of the Board of Governors, including the Chair and Vice Chair. Macro variables include inflation, real wage growth, and GDP growth.

A.46 Dummy for 1993-Q4

Table 14: Results for H2, with dummy for 1993-Q4

	1	2	3	4	5
Unemployment	−0.072	−0.065	−0.048	−0.055	−0.044
	(0.042)	(0.043)	(0.045)	(0.044)	(0.043)
Unemployment^2	0.043†	0.045‡	0.042†	0.041†	0.038†
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
Forum (FOMC=1)	0.131	0.139	0.146	0.141	0.119
	(0.109)	(0.109)	(0.109)	(0.108)	(0.107)
Unemp*Forum	−0.154‡	−0.153‡	−0.150‡	−0.151‡	−0.158‡
	(0.050)	(0.049)	(0.049)	(0.049)	(0.049)
Unemp^2*Forum	−0.001	−0.002	−0.003	−0.002	−0.002
	(0.023)	(0.022)	(0.022)	(0.022)	(0.022)
Real wage growth				0.024	0.028
				(0.029)	(0.029)
Nominal wage growth			0.050		
			(0.030)		
Inflation		0.036		0.057	0.067†
		(0.019)		(0.032)	(0.032)
Real GDP growth					0.034‡
					(0.013)
1993-Q4	0.621‡‡	0.683‡‡	0.644‡‡	0.692‡‡	0.683‡‡
	(0.122)	(0.128)	(0.123)	(0.128)	(0.127)
Quarterly word count (log)	0.392‡‡	0.397‡‡	0.395‡‡	0.396‡‡	0.381‡‡
	(0.098)	(0.098)	(0.098)	(0.098)	(0.097)
Num.Obs.	332	332	332	332	332
AIC	2320.4	2319.4	2319.7	2320.8	2315.7
BIC	2373.7	2376.5	2376.8	2381.7	2380.3
Chair dummies	✓	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Dummy for 1993-Q4 corresponds to the quarter when it became known to Fed officials that FOMC trancripts would be made public with a five-year lag. Note that effects on *Unemp*Forum* remain significant in all specifications. While the dummy on 1993-Q4 is significant, we tested other models with dummies for 1991-Q4 and 1995-Q4 as placebos and got similar significance, indicating that this significance instead more likely reflects the longer-term shifts captured in Section 5.5.

A.5 Robustness checks for models in Section 5.4

A.51 Tealbook data

Table 15: Results on H3, a partisan fear of full employment, 1978-2019

	1	2	3	4
Unemployment (GB)	−0.165‡‡	−0.165‡‡	−0.113	−0.082
	(0.038)	(0.038)	(0.061)	(0.060)
Unemployment^2 (GB)	0.062‡‡	0.061‡‡	0.074‡‡	0.070‡‡
	(0.015)	(0.015)	(0.020)	(0.020)
Forum (FOMC=1)	0.162	−0.088	0.004	−0.023
	(0.102)	(0.118)	(0.135)	(0.131)
Unemp*Forum	−0.184‡‡	−0.203‡‡	−0.317‡‡	−0.335‡‡
	(0.045)	(0.044)	(0.062)	(0.059)
President (Dem=1)	0.423‡‡	0.168	0.223	0.247
	(0.073)	(0.095)	(0.131)	(0.129)
Pres*Forum		0.514‡‡	0.358†	0.384†
		(0.128)	(0.179)	(0.173)
Pres*Unemp			−0.120	−0.144
			(0.082)	(0.082)
Unemp*Forum*Pres			0.220†	0.238‡
			(0.093)	(0.088)
Num.Obs.	332	332	334	332
AIC	2281.1	2267.4	2301.1	2258.1
BIC	2345.8	2335.9	2373.5	2341.8
Chair dummies	✓	✓	✓	✓
Macro controls	✓	✓		✓
Log(total words)	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: "GB" stands for Greenbook. Greenbook variables are not lagged; wage variables are lagged by one quarter to account for release dates. Macro controls include real wage growth, inflation, and real GDP growth. All interaction terms including *Unemployment^2* are included in model but do not have significant effects and are ommitted from the table for simplicity.

.

A.52 Unemployment gap

Table 16: Results on H3, a partisan fear of full employment, 1978-2019

	1	2	3	4
UGAP	−0.201†	−0.229‡	0.076	0.084
	(0.081)	(0.080)	(0.116)	(0.120)
UGAP^2	0.036	0.046†	0.028	0.028
	(0.022)	(0.021)	(0.033)	(0.034)
Forum (FOMC=1)	0.222†	0.012	0.157	0.152
	(0.101)	(0.116)	(0.120)	(0.119)
UGAP*Forum	−0.152	−0.125	−0.326†	−0.339†
	(0.106)	(0.104)	(0.148)	(0.147)
President (Dem=1)	0.368‡‡	0.114	0.391‡‡	0.405‡‡
	(0.082)	(0.107)	(0.119)	(0.118)
Pres*Forum		0.503‡‡	0.241	0.218
		(0.144)	(0.157)	(0.155)
Pres*UGAP			−0.571‡‡	−0.535‡‡
			(0.160)	(0.161)
UGAP*Forum*Pres			0.470†	0.473†
			(0.206)	(0.203)
Num.Obs.	332	332	332	332
AIC	2342.1	2332.2	2317.7	2314.4
BIC	2406.8	2400.7	2390.0	2398.1
Chair dummies	✓	✓	✓	✓
Macro controls	✓	✓		✓
Log(total words)	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: Macro controls include real wage growth, inflation, and real GDP growth. All interaction terms including *Unemployment^2* are included in model but do not have significant effects and are omitted from the table for simplicity.

A.53 Congressional regressions

Table 17: Results on H3, a partisan fear of full employment, 1978-2019

	1	2	3	4
Unemployment	−0.110†	−0.108†	−0.041	−0.018
	(0.048)	(0.048)	(0.062)	(0.068)
Unemployment^2	0.049‡	0.048‡	0.055‡	0.055‡
	(0.017)	(0.017)	(0.021)	(0.021)
Forum (FOMC=1)	0.123	−0.131	−0.109	−0.119
	(0.105)	(0.123)	(0.139)	(0.137)
Unemp*Forum	−0.160‡‡	−0.185‡‡	−0.289‡‡	−0.300‡‡
	(0.048)	(0.047)	(0.063)	(0.062)
President (Dem=1)	0.339‡‡	0.084	0.093	0.149
	(0.079)	(0.102)	(0.138)	(0.138)
Pres*Forum		0.510‡‡	0.488‡	0.466‡
		(0.133)	(0.182)	(0.180)
Pres*Unemp			−0.168	−0.156
			(0.086)	(0.088)
Unemp*Forum*Pres			0.235†	0.241†
			(0.097)	(0.096)
House (Dem=1)	−0.265†	−0.269†	−0.297†	−0.230
	(0.121)	(0.119)	(0.118)	(0.121)
Senate (Dem=1)	−0.021	−0.020	0.030	−0.065
	(0.142)	(0.139)	(0.129)	(0.144)
Num.Obs.	332	332	332	332
AIC	2310.9	2298.8	2299.8	2294.9
BIC	2383.2	2374.9	2379.7	2386.2
Chair dummies	✓	✓	✓	✓
Macro controls	✓	✓		✓
Log(total words)	✓	✓	✓	✓

† p < 0.05, ‡ p < 0.01, ‡‡ p < 0.001

Note: Macro controls include real wage growth, inflation, and real GDP growth. All interaction terms including *Unemployment^2* are included in model but do not have significant effects and are ommitted from the table for simplicity.

Figure 21: Controlling party of US Congress and president (quarterly)

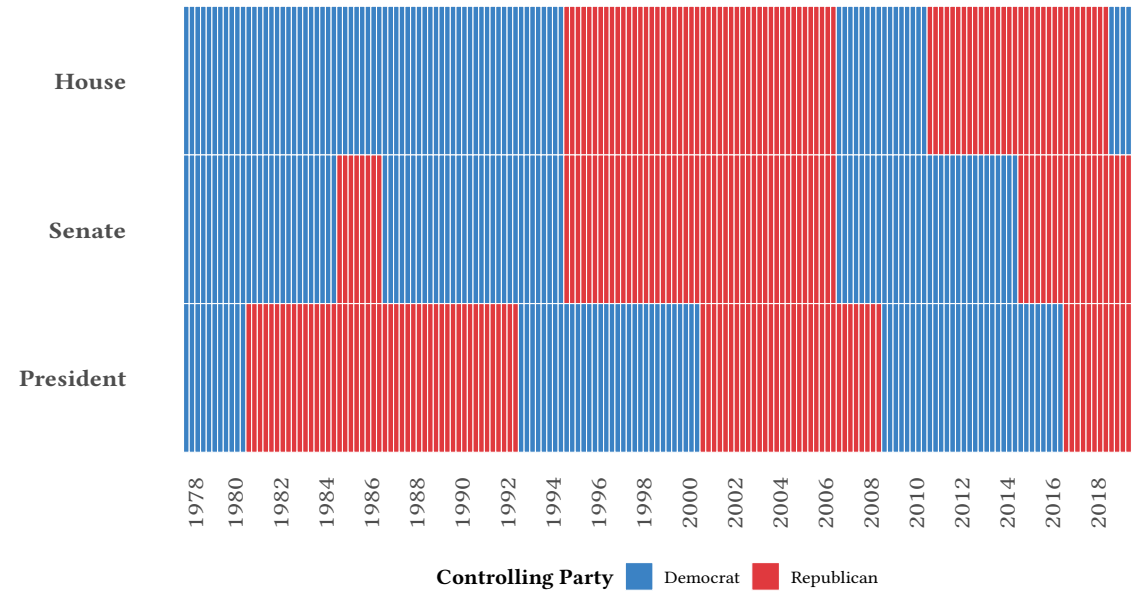

Figure 22: Rolling 15-year regressions, partisan variables

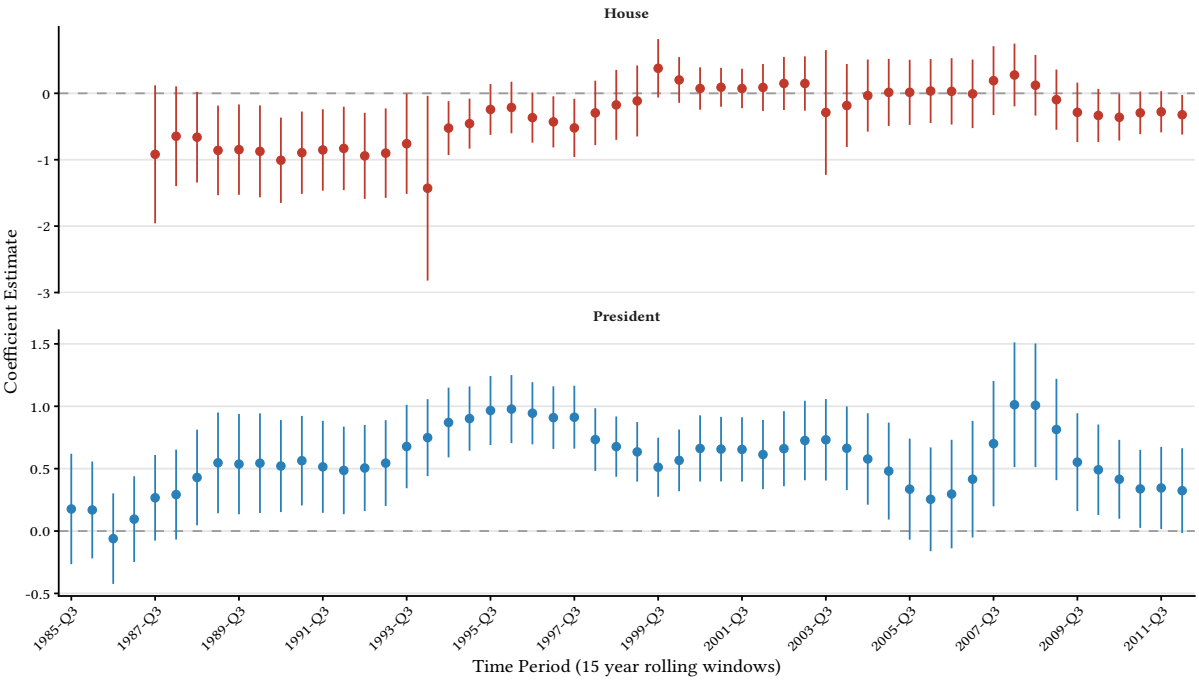

<>

A.6 Further robustness from Section 5.5

Figure 23: Leave-one-out analysis, model 5 in Table 2