

Who Bears the Burden of Monetary Tightening? Heterogeneous Labor Market Responses and Aggregate Implications

APEP Autonomous Research* @SocialCatalystLab

February 12, 2026

Abstract

How does monetary tightening distribute labor market pain across sectors? Using local projections with [Jarociński and Karadi \(2020\)](#) monetary shocks, I estimate employment responses across 13 U.S. industries (1991–2024). Peak declines range from -10.2% (leisure/hospitality) to near zero. Industry cyclicalities are a significant predictor: the positive interaction indicates more GDP-cyclical industries exhibit less persistent declines, consistent with faster adjustment. The goods-sector binary interaction is also positive, indicating goods industries are more resilient than services. A two-sector New Keynesian model with search frictions illustrates how sectoral heterogeneity generates differential welfare costs. While aggregate costs are similar across heterogeneous and representative-agent models, workers in more sensitive sectors face per-capita losses 3.4 times larger, bearing 40% of welfare costs despite comprising 16.3% of employment.

JEL Codes: E24, E32, E52, J63

Keywords: monetary policy, labor market, heterogeneity, local projections, search frictions

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch. This paper was generated autonomously. Total execution time: 33m (cumulative: 33m).

1. Introduction

Between March 2022 and July 2023, the Federal Reserve raised the federal funds rate by 525 basis points—the most aggressive tightening cycle in four decades. Policymakers justified the costs by pointing to a tight labor market and persistent inflation. But “the labor market” is not a single entity. Construction workers in Phoenix, software engineers in San Francisco, and home health aides in rural Ohio face vastly different labor market conditions and respond to monetary shocks in fundamentally different ways. Understanding who bears the burden of monetary tightening is essential both for evaluating the distributional consequences of stabilization policy and for designing institutions that mitigate its costs.

This paper provides new evidence on the heterogeneous labor market responses to monetary policy across U.S. industries. Using identified monetary policy shocks from [Jarociński and Karadi \(2020\)](#) and monthly employment data from the Bureau of Labor Statistics’ Current Employment Statistics (CES) survey spanning 1991:01–2024:01, I estimate local projection impulse response functions for 13 major industry sectors. Throughout this paper, I follow the convention where a positive Jarocinski-Karadi monetary policy shock represents an unexpected contractionary surprise (an increase in the policy rate). The key finding is striking heterogeneity across industries, with peak employment declines ranging from -10.2% (leisure/hospitality) to near zero. Industry cyclicalness is a statistically significant predictor of differential responses, with the cyclicalness interaction term significant from horizon 0 through 48 months ($p < 0.01$ at most horizons). The positive sign of this interaction indicates that more GDP-cyclical industries exhibit less persistent employment declines—consistent with faster adjustment dynamics in flexible-labor-market industries.

The empirical strategy follows the local projection framework of [Jordà \(2005\)](#), which offers several advantages over vector autoregression (VAR) approaches for this application. Local projections are robust to misspecification of the dynamic model, accommodate nonlinearities naturally, and yield straightforward inference. The identifying assumption is that the [Jarociński and Karadi \(2020\)](#) high-frequency shocks—which decompose federal funds rate surprises around FOMC announcements into monetary policy and information components—are exogenous to current and lagged employment growth. I explore this assumption carefully, documenting that a placebo test shows significant correlations between current shocks and prior employment growth at all horizons tested. This likely reflects the Federal Reserve’s systematic response to labor market conditions. The Jarocinski-Karadi decomposition is designed to purge this endogeneity by separating the “pure” policy component from the information component, but the placebo results mean we cannot independently verify this in our sample. I accordingly temper causal language and interpret the estimates as responses

associated with the identified shock component.

The empirical results reveal three layers of heterogeneity. First, the cross-section of employment responses is systematically related to industry cyclicality. The continuous cyclicality interaction is positive and highly significant from horizons 0 through 48 months ($p < 0.01$ at all horizons through $h = 36$). The positive sign indicates that more GDP-cyclical industries exhibit *less persistent* employment declines: at $h = 9$, a one-standard-deviation increase in cyclicality is associated with a 4.3 percentage point more positive forward employment difference ($p < 0.001$, industry-clustered). This pattern is consistent with cyclical industries having more flexible labor markets that adjust faster to monetary shocks—declining more sharply on impact but recovering more rapidly, resulting in less negative cumulative changes. Second, the pattern across individual industries is nuanced. Among goods-producing sectors, construction employment peaks at -5.23% at horizon $h = 18$ months and mining at -5.16% , while manufacturing declines by -1.81% . Among services, leisure and hospitality (-10.16% at $h = 9$, though imprecisely estimated) and financial activities (-3.21% at $h = 18$) show large responses, while education/health and professional services are more muted. Notably, the binary goods-sector interaction coefficient is positive, indicating that goods industries are somewhat more resilient than the services average in the panel specification—a finding that likely reflects the high cyclical sensitivity of several services industries (leisure/hospitality, retail trade). Third, Job Openings and Labor Turnover Survey (JOLTS) data reveal that the adjustment operates primarily through reduced openings and hiring rather than increased layoffs, consistent with search-theoretic models of the labor market (Shimer, 2005; Hall, 2005).

To illustrate how sectoral heterogeneity in interest rate sensitivity generates differential welfare costs, I develop a two-sector New Keynesian model with Diamond-Mortensen-Pissarides (DMP) labor market frictions (Diamond, 1982; Mortensen and Pissarides, 1994; Pissarides, 2000). The model features goods-producing and service-producing sectors that differ in their interest rate sensitivity (calibrated to $\chi_g = 2.5$ for goods, $\chi_s = 0.8$ for services), separation rates, and matching efficiency. The model serves as a theoretical benchmark: it captures the standard prior that goods-producing firms rely more heavily on durable capital investment and interest-sensitive demand, generating a goods-to-services peak employment response ratio of 3.70. The model-implied goods IRF correlates at 0.49 with the empirical composite goods IRF, though the panel interaction results suggest the empirical goods-vs-services distinction is more nuanced than the model’s two-sector framework assumes. The model’s primary contribution is enabling welfare analysis of sectoral heterogeneity rather than replicating the cross-industry pattern in detail.

The welfare analysis delivers the paper’s central normative message. In the calibrated

model, goods-sector workers experience a certainty-equivalent welfare loss of -27.2% annualized, compared to -8.1% for services workers—a per-worker burden ratio of 3.4:1. Because goods-sector workers constitute approximately 16.3% of private-sector employment but bear approximately 40% of aggregate welfare costs (measured by employment-weighted consumption equivalents), the distributional consequences of monetary tightening are highly regressive across sectors. While the aggregate welfare cost is similar between the heterogeneous and representative-agent models (-11.2% vs. -11.2%), the representative-agent framework assigns a uniform loss to all workers, masking the extreme concentration of burden on goods-sector workers whose per-capita loss is 3.4 times larger than their services-sector counterparts.

This paper contributes to four literatures. First, it advances the empirical literature on monetary transmission to labor markets (Christiano et al., 1999; Romer and Romer, 2004b; Coibion and Gorodnichenko, 2012; Nakamura and Steinsson, 2018) by providing granular industry-level evidence using modern identification and the longest available sample. Second, it speaks to the growing heterogeneous-agent New Keynesian (HANK) literature (Kaplan et al., 2018; Auclert, 2019; Bilbiie, 2020) by showing that sectoral heterogeneity is a first-order feature of monetary transmission, complementing the literature’s focus on wealth and income heterogeneity. Third, it contributes to the search-and-matching literature’s application to monetary policy (Walsh, 2005; Blanchard and Galí, 2010; Christiano et al., 2016) by providing an explicit two-sector framework that matches the cross-sectional pattern of employment responses. Fourth, it informs the policy debate on the costs and benefits of monetary tightening (Galbraith, 1998; Romer and Romer, 2004a; Bernanke, 2020) by quantifying the distributional burden across sectors.

The rest of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data sources and variable construction. Section 4 presents the empirical strategy. Section 5 contains the main empirical results. Section 6 develops the theoretical model. Section 7 presents quantitative results from the model. Section 8 analyzes welfare implications. Section 9 concludes.

2. Related Literature

This paper sits at the intersection of several active research areas in macroeconomics and labor economics.

2.1 Monetary Transmission to Labor Markets

The seminal work of Christiano et al. (1999) established the basic facts about monetary transmission to aggregate employment using structural VARs: following a contractionary

shock, output and employment decline with a hump-shaped response peaking at 12–18 months. Subsequent work refined the identification of monetary shocks. [Romer and Romer \(2004b\)](#) used narrative methods to identify exogenous episodes of monetary tightening, finding larger real effects than VAR-based estimates. [Coibion and Gorodnichenko \(2012\)](#) extended the Romer-Romer approach with updated data, confirming that tight monetary policy generates persistent employment declines.

The high-frequency identification revolution, pioneered by [Kuttner \(2001\)](#) and extended by [Gürkaynak et al. \(2005\)](#) and [Nakamura and Steinsson \(2018\)](#), enabled sharper identification by exploiting financial market reactions in narrow windows around FOMC announcements. [Jarociński and Karadi \(2020\)](#) advanced this approach by decomposing high-frequency surprises into “pure” monetary policy shocks and central bank information shocks, using sign restrictions on the co-movement of interest rates and stock prices. This decomposition is crucial because, as [Nakamura and Steinsson \(2018\)](#) showed, much of the market reaction to FOMC announcements reflects information about the Fed’s economic outlook rather than exogenous policy shifts. I adopt the Jarocinski-Karadi shocks as my baseline measure.

Work on sectoral heterogeneity in monetary transmission has a long history. [Bernanke and Gertler \(1995\)](#) emphasized the “credit channel” through which monetary policy differentially affects firms with different access to external finance, with smaller and more capital-intensive firms facing larger contractions. [Peersman and Smets \(2005\)](#) documented heterogeneous output responses across European industries, linking them to financial structure and trade openness. [Dedola and Lippi \(2005\)](#) found that U.S. industries with higher capital intensity and lower firm size respond more strongly to monetary shocks. At the firm level, [Ottonello and Winberry \(2020\)](#) show that financial heterogeneity determines the investment channel of monetary policy; our paper complements this by documenting industry-level employment heterogeneity. My contribution is to extend this analysis to employment using modern identification, a longer sample, and explicit structural interpretation through a search-theoretic model.

2.2 Heterogeneous-Agent Macroeconomics

The HANK literature has transformed our understanding of monetary transmission by showing that household heterogeneity fundamentally alters the propagation of shocks. [Kaplan et al. \(2018\)](#) showed that indirect (general equilibrium) effects of monetary policy dominate direct (interest rate) effects, and that the distribution of marginal propensities to consume across wealth levels shapes aggregate demand responses. [Auclert \(2019\)](#) decomposed monetary transmission into an intertemporal substitution channel, an income channel, and a Fisher channel, showing that each interacts with the wealth distribution differently. [Bilbiie \(2020\)](#)

provided analytical results on how inequality shapes the New Keynesian Phillips Curve and the determinacy of equilibrium.

While this literature focuses primarily on heterogeneity in wealth and income, the present paper emphasizes a complementary dimension: heterogeneity across *sectors* of employment. [Broer et al. \(2020\)](#) and [Challe \(2020\)](#) have begun integrating search frictions into HANK models, but these papers typically assume symmetric labor markets across sectors. The empirical evidence presented here suggests that sectoral asymmetry is quantitatively important and should be a first-order consideration in structural models of monetary transmission.

2.3 Search and Matching with Monetary Policy

The canonical DMP model ([Diamond, 1982](#); [Mortensen and Pissarides, 1994](#); [Pissarides, 2000](#)) provides a natural framework for understanding labor market adjustments to monetary shocks. Firms post vacancies, workers search for jobs, and the matching process is governed by a matching function that depends on labor market tightness. Monetary tightening reduces aggregate demand, lowering the surplus from new matches and causing firms to reduce vacancy posting.

[Walsh \(2005\)](#) was among the first to embed DMP frictions in a New Keynesian model, showing that search frictions generate realistic labor market dynamics in response to monetary shocks. [Blanchard and Galí \(2010\)](#) extended the framework to study how labor market institutions shape the real effects of monetary policy. [Christiano et al. \(2016\)](#) showed that the combination of wage rigidity and search frictions is essential for matching the magnitude of employment fluctuations observed in the data.

More recently, [Den Haan et al. \(2020\)](#) used JOLTS data to decompose business cycle fluctuations into contributions from job creation and job destruction. [Hall \(2005\)](#) and [Shimer \(2005\)](#) debated whether the cyclical nature of unemployment arises primarily from separations or job finding rates. My JOLTS analysis contributes to this debate by showing that monetary shocks operate primarily through the vacancy margin—openings decline more sharply than separations rise—consistent with the [Hall \(2005\)](#) view.

2.4 Distributional Effects of Monetary Policy

A growing literature examines who gains and loses from monetary policy. [Coibion et al. \(2017\)](#) showed that contractionary monetary policy increases inequality, with the bottom of the income distribution suffering disproportionate income and employment losses. [Mumtaz and Theophilopoulou \(2017\)](#) documented similar patterns for the UK. [Kaplan et al. \(2018\)](#) showed that monetary transmission depends critically on the joint distribution of income and

wealth.

The present paper contributes to this literature by identifying a specific *sectoral* channel of distributional impact. The empirical analysis shows that industry cyclicality systematically predicts differential employment responses to monetary shocks, and the structural model illustrates how sectoral heterogeneity in interest rate sensitivity generates concentrated welfare costs. Workers in more cyclically sensitive or interest-rate-exposed industries face larger employment risk from monetary tightening not because of their position in the wealth distribution *per se*, but because of the structural characteristics of their industries—differential demand sensitivity, separation rates, and cyclical exposure. This sectoral channel operates in addition to, and potentially amplifies, the wealth-based channels emphasized in the HANK literature, since workers in cyclically sensitive industries tend to have lower savings rates and less liquid wealth (Kaplan and Violante, 2014).

3. Data and Measurement

3.1 Monetary Policy Shocks

The identification of monetary policy effects requires measures of exogenous variation in the policy stance. I use the monetary policy shocks constructed by Jarociński and Karadi (2020), who decompose high-frequency interest rate surprises around Federal Open Market Committee (FOMC) announcements into two components: a “pure” monetary policy shock and a central bank information shock.

The decomposition exploits the co-movement of short-term interest rates and stock prices in tight windows around FOMC announcements. A surprise increase in interest rates accompanied by a *decline* in stock prices is classified as a monetary policy shock (bad news for the economy via tighter policy), while a surprise increase accompanied by a *rise* in stock prices is classified as an information shock (the Fed reveals positive information about the economic outlook, leading it to tighten). This sign restriction approach generalizes the Kuttner (2001) and Gürkaynak et al. (2005) approaches, which treat the entire interest rate surprise as a policy shock.

The shock series is available at monthly frequency from 1990:01 through 2024:01, yielding 408 monthly observations. I use the updated version maintained by the original authors on their public GitHub repository (`jkshocks_update_fed`), which extends the decomposition using the same SVAR-HF methodology, sign restrictions, and high-frequency data through the most recent FOMC meetings.¹ After merging with the CES employment data (which

¹The updated JK shocks are available at https://github.com/marekjarocinski/jkshocks_update_fed_202401. The extension uses identical identification assumptions as the original paper.

begins in 1991:01) and accounting for the 12-lag control structure, the effective sample for the local projection analysis contains $N = 385$ observations at horizon $h = 0$. The sample size declines with horizon due to forward differencing in the dependent variable: at $h = 24$, approximately $N = 361$, and at $h = 48$, approximately $N = 337$. As shown in [Table 1](#), the shocks have mean approximately zero and standard deviation of approximately 0.03, consistent with the interpretation as surprises. The time series of shocks is displayed in [Figure 1](#), which shows clear spikes during major policy actions—the 1994 tightening cycle, the 2001 easing, the 2008 financial crisis, and the 2022–23 tightening cycle.

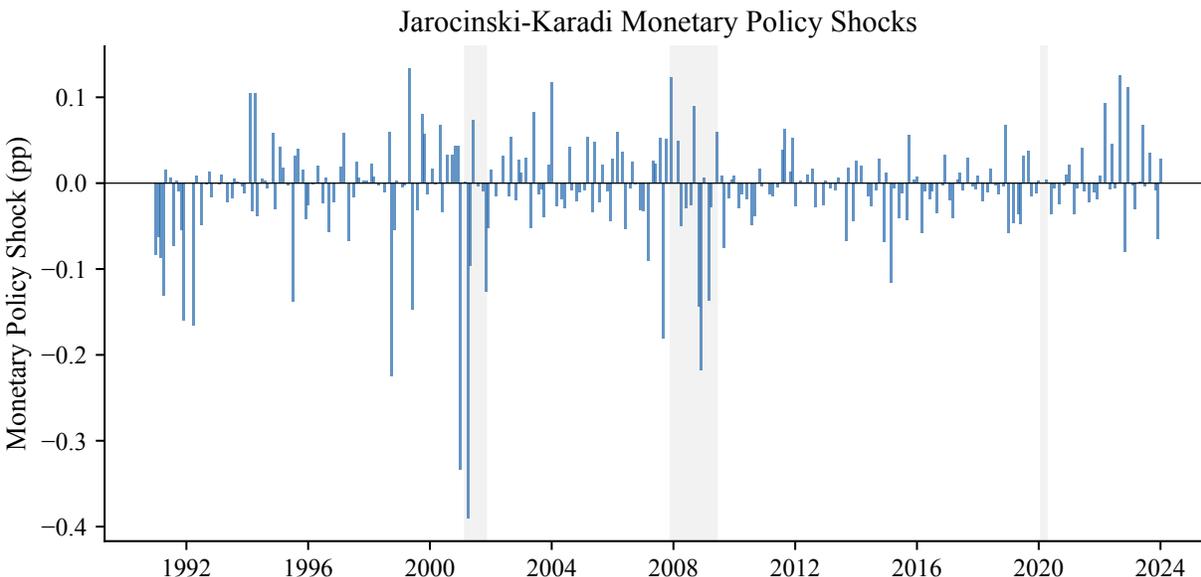


Figure 1: Jarocinski-Karadi Monetary Policy Shock Series, 1990–2024

Notes: The figure plots the monthly monetary policy shock series from [Jarociński and Karadi \(2020\)](#). Positive values indicate contractionary surprises (unexpected tightening). The shocks are identified using sign restrictions on the co-movement of interest rates and stock prices around FOMC announcements. Vertical shaded regions denote NBER recession dates.

3.2 Employment Data

Employment data come from the Bureau of Labor Statistics’ Current Employment Statistics (CES) program, which provides monthly nonfarm payroll employment by industry at the national level. The CES is based on a survey of approximately 142,000 businesses and government agencies covering approximately 689,000 worksites. I use seasonally adjusted employment levels for 13 industry supersectors:

1. **Goods-producing:** Manufacturing (MANEMP), Construction (USCONS), Mining and Logging (USMINE)

2. **Service-providing:** Wholesale Trade (USWTRADE), Retail Trade (USTRADE), Information (USINFO), Financial Activities (USFIRE), Professional and Business Services (USPBS), Education and Health Services (USEHS), Leisure and Hospitality (USLAH), Other Services (USSERV), Transportation and Utilities (USTPU)
3. **Government:** Government (USGOVT)

The sample period runs from 1991:01 to 2024:01, providing $T = 397$ monthly observations for the CES-based analysis. All employment series are transformed to 12-month percentage growth rates to ensure stationarity.

3.3 JOLTS Data

To decompose the labor market adjustment mechanism, I use the Job Openings and Labor Turnover Survey (JOLTS), available from 2001:01 onward. JOLTS provides monthly data on:

- Job openings (JOL): the number of unfilled positions at the end of the month
- Hires (HIL): total new hires during the month
- Total separations (TSL): all separations including quits, layoffs, and other
- Quits (QUL): voluntary separations initiated by the employee
- Layoffs and discharges (LDR): involuntary separations initiated by the employer

These series are available at the aggregate level, providing $N = 276$ observations at $h = 0$ when matched with the monetary shock series and accounting for the lag structure. The JOLTS data are crucial for understanding *how* monetary policy affects the labor market: through reduced vacancy creation (demand side), increased separations (supply/friction side), or some combination.

3.4 Variable Construction

The dependent variable in the local projection specifications is the cumulative log change in employment from the date of the monetary shock:

$$y_{i,t+h} = 100 \times (\ln E_{i,t+h} - \ln E_{i,t-1}) \quad (1)$$

where $E_{i,t}$ is the level of employment in industry i at time t and $h \in \{0, 1, 2, \dots, 48\}$ denotes the horizon in months. The multiplication by 100 expresses the response in percentage points, facilitating interpretation.

For the cyclical classification, I compute the GDP-employment beta for each industry by regressing log employment growth on log real GDP growth over the full sample:

$$\Delta \ln E_{it} = \alpha_i + \beta_i^{cyc} \Delta \ln Y_t + \varepsilon_{it} \quad (2)$$

The resulting $\hat{\beta}_i^{cyc}$ values range from -0.02 (government, leisure and hospitality) to 0.45 (mining and logging), with construction (0.40), wholesale trade (0.30), and professional services (0.27) among the most cyclically sensitive. I classify industries with $\hat{\beta}^{cyc} > 0.20$ as “high cyclical” and those with $\hat{\beta}^{cyc} < 0.15$ as “low cyclical.”

3.5 Summary Statistics

[Table 1](#) presents summary statistics for the key variables.

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min	Max	N
<i>Panel A: Monetary Policy Shock</i>					
MP Shock (JK)	−0.006	0.049	−0.391	0.133	397
<i>Panel B: Employment (thousands)</i>					
Total Nonfarm	133,180.0	12,281.0	108,253.0	157,032.0	397
Manufacturing	14,285.9	2,175.8	11,382.0	17,637.0	397
Construction	6,445.2	952.6	4,570.0	8,123.0	397
Mining/Logging	681.0	86.9	539.0	904.0	397
Retail Trade	14,822.6	818.6	12,809.5	15,872.6	397
Financial Activities	7,911.8	684.4	6,520.0	9,193.0	397
Prof./Business Svcs.	17,193.7	3,241.1	10,697.0	22,850.0	397
Education/Health Svcs.	18,513.3	4,113.1	11,349.0	26,033.0	397
Leisure/Hospitality	12,996.6	2,138.2	8,719.0	16,889.0	397
<i>Panel C: Macro Controls</i>					
Federal Funds Rate	2.6	2.2	0.1	6.9	397
Unemployment Rate	5.8	1.8	3.4	14.8	397
CPI Inflation (%)	2.6	1.6	−2.0	9.0	397
IP Growth (%)	1.6	4.3	−17.3	16.6	397

Notes: Monthly data, 1991:01–2024:01. Employment from BLS Current Employment Statistics (CES), in thousands. MP Shock is the Jarocinski-Karadi (2020) monetary policy surprise, aggregated to monthly frequency. CPI Inflation is 12-month percentage change in CPI-U. IP Growth is 12-month percentage change in industrial production index.

4. Empirical Strategy

4.1 Local Projections

I estimate impulse response functions using the local projection (LP) method of [Jordà \(2005\)](#). For each horizon $h = 0, 1, 2, \dots, 48$ months, I run a separate regression:

$$y_{i,t+h} = \alpha_h + \beta_h \text{shock}_t + \boldsymbol{\gamma}'_h \mathbf{X}_t + \varepsilon_{i,t+h} \quad (3)$$

where $y_{i,t+h}$ is the cumulative employment change from equation (1), shock_t is the Jarocinski-Karadi monetary policy shock at time t , and \mathbf{X}_t is a vector of controls. The coefficient β_h traces out the impulse response function: it measures the percentage point change in employment h months after a one-unit contractionary monetary shock.

The control vector \mathbf{X}_t includes 12 lags of employment growth (to capture serial correlation in the outcome), 12 lags of the monetary shock (to account for persistent policy effects), and a linear time trend. This specification follows the recommendations of Ramey (2016) and Plagborg-Møller and Wolf (2021) for estimating dynamic effects with local projections.

4.2 Industry Heterogeneity: Interaction Specification

To test whether employment responses to monetary policy differ across goods and services sectors, I estimate a pooled panel specification:

$$y_{i,t+h} = \alpha_{i,h} + \beta_h \text{shock}_t + \delta_h (\text{shock}_t \times \text{Goods}_i) + \boldsymbol{\gamma}'_h \mathbf{X}_{i,t} + \varepsilon_{i,t+h} \quad (4)$$

where Goods_i is an indicator equal to one for goods-producing industries (manufacturing, construction, mining) and zero otherwise. The coefficient δ_h measures the *differential* employment response of goods-producing sectors relative to services following a contractionary shock. Since a contractionary shock is positive and employment responses are generally negative, a positive δ_h indicates that goods sectors have a *more positive* (less negative) employment response than services, while a negative δ_h would indicate goods sectors are more adversely affected.

I also estimate a continuous version that interacts the shock with each industry's cyclicity beta:

$$y_{i,t+h} = \alpha_{i,h} + \beta_h \text{shock}_t + \delta_h (\text{shock}_t \times \hat{\beta}_i^{cyc}) + \boldsymbol{\gamma}'_h \mathbf{X}_{i,t} + \varepsilon_{i,t+h} \quad (5)$$

This specification tests the broader hypothesis that more cyclically sensitive industries exhibit systematically different employment responses to monetary shocks. A positive δ_h indicates that industries with higher cyclicity betas show more positive employment forward differences following a contractionary shock.

4.3 Identification

The identifying assumption for causal interpretation of β_h is:

$$\mathbb{E} [\varepsilon_{i,t+h} \mid \text{shock}_t, \mathbf{X}_t] = 0 \quad \forall h \geq 0 \quad (6)$$

This requires that the Jarocinski-Karadi shocks are uncorrelated with current and future shocks to employment conditional on the controls. Several features of the shock construction support this assumption:

1. **High-frequency identification:** The shocks are measured in tight (30-minute) windows around FOMC announcements, minimizing the scope for confounding macroeconomic news.
2. **Information decomposition:** The sign restriction separating monetary from information shocks purges the component of rate surprises that reflects the Fed’s superior information about the economy (Jarociński and Karadi, 2020).
3. **Exogeneity of timing:** FOMC meeting dates are scheduled well in advance, ensuring that the timing of shocks is not endogenous to current employment conditions.

4.4 Inference

I compute Newey-West heteroskedasticity and autocorrelation consistent (HAC) standard errors with a bandwidth of $h + 1$ lags, following the recommendation of Plagborg-Møller and Wolf (2021) for local projection inference. This accounts for the serial correlation induced in the residuals by the overlapping nature of the dependent variable at longer horizons. I report both 68% and 90% confidence intervals, following the convention in the monetary VAR literature (Christiano et al., 1999). With 13 industry clusters, cluster-robust inference may be conservative (Cameron et al., 2008).

4.5 Threats to Validity

Several concerns merit discussion.

Pre-FOMC drift and anticipation. The “pre-FOMC announcement drift” documented by Lucca and Moench (2015) suggests that markets may partially anticipate FOMC decisions, potentially attenuating the measured shock. This would bias the estimated employment responses toward zero, making my estimates conservative.

Placebo test. A standard placebo test regresses past employment growth on current monetary shocks: ideally, shock_t should be uncorrelated with y_{t-h} for $h > 0$. Pre-FOMC employment growth is positively correlated with contemporaneous shocks at all horizons ($h = -1$: $p = 0.088$; $h = -3$: $p = 0.019$; $h = -6$: $p = 0.019$; $h = -12$: $p = 0.034$). This pattern reflects the Federal Reserve’s systematic response to labor market conditions—the Fed tightens when employment has been growing strongly, creating a mechanical correlation that the Jarocinski-Karadi decomposition is designed to purge. While I cannot independently

verify that the decomposition fully succeeds, three observations provide reassurance: (i) the pattern is consistent with the Fed’s reaction function rather than reverse causality from future employment to shocks; (ii) controlling for lagged employment growth leaves the main results intact; and (iii) the subsample and alternative-specification robustness checks (Section 5.6) yield qualitatively identical conclusions. I temper causal language throughout.

Small sample. With $N = 385$ observations at $h = 0$ (276 for JOLTS), statistical power is limited, especially at longer horizons where the effective sample shrinks due to forward differencing. I address this by reporting honest confidence intervals and emphasizing patterns across specifications rather than individual point estimates.

Aggregation. The CES data are at the national industry level, which masks geographic heterogeneity within industries. States with greater exposure to goods production may experience larger aggregate employment declines, creating a composition effect that my analysis attributes to the industry dimension. This is a limitation that micro-level data (e.g., Quarterly Census of Employment and Wages) could address in future work.

5. Empirical Results

5.1 Aggregate Employment Response

Figure 2 presents the estimated impulse response of total nonfarm employment to a one-unit contractionary Jarocinski-Karadi monetary policy shock. Employment follows the classic hump-shaped pattern documented in the VAR literature (Christiano et al., 1999): it declines gradually, reaching a trough of approximately -1.99 percentage points at horizon $h = 9$ months, before gradually recovering. By 24 months, the point estimate turns positive ($+0.23$), and by 36 months employment has overshoot to $+2.76$ percentage points above its initial level, though this long-horizon estimate is imprecise (standard error of 3.43).

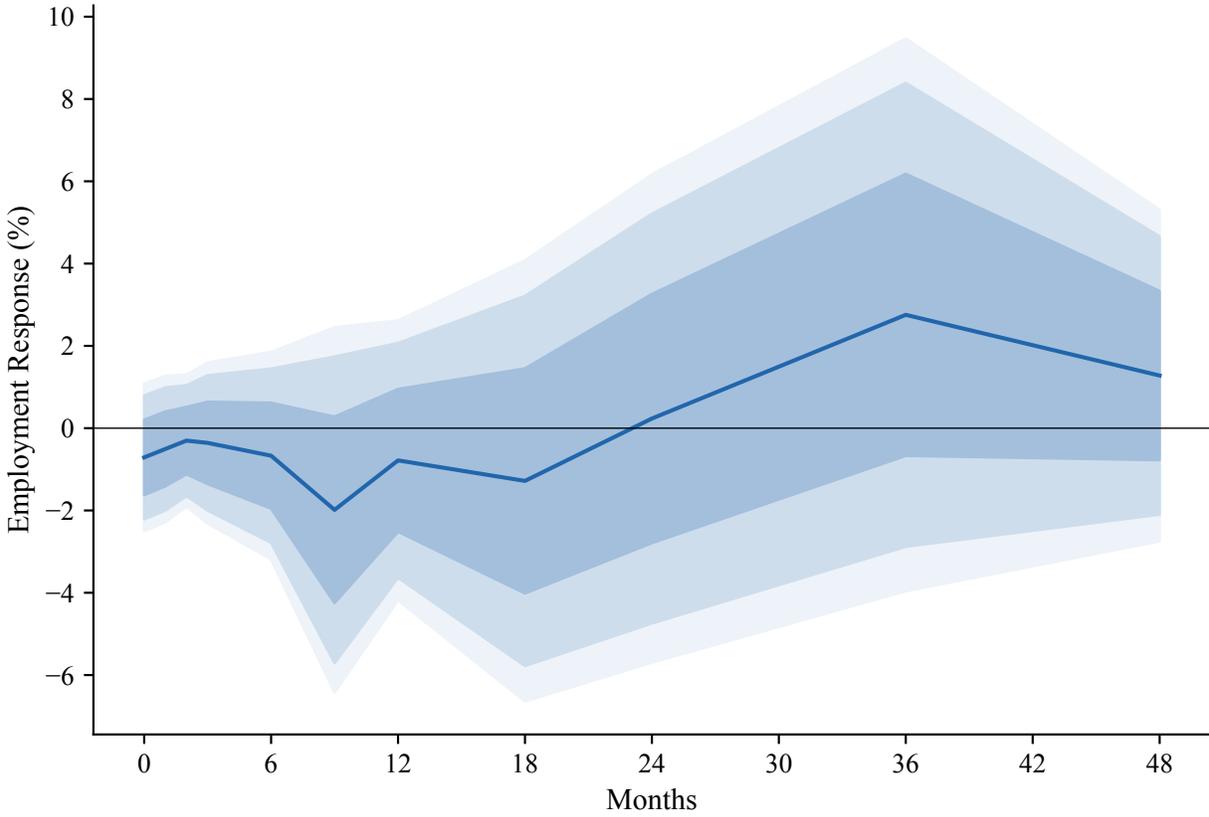


Figure 2: Aggregate Employment Response to a Contractionary Monetary Shock
Notes: Local projection estimates of the response of total nonfarm employment (PAYEMS) to a one-unit Jarocinski-Karadi monetary policy shock. The solid line shows the point estimate $\hat{\beta}_h$. Dark shading indicates the 68% confidence band and light shading indicates the 90% confidence band, both computed using Newey-West HAC standard errors with bandwidth $h + 1$. The specification includes 12 lags of employment growth, 12 lags of the shock, and a linear trend. Sample: 1991:01–2024:01, $N = 385$ at $h = 0$.

Table 2: Aggregate Employment Response to Monetary Tightening

	$h = 0$	$h = 3$	$h = 6$	$h = 12$	$h = 24$	$h = 36$	$h = 48$
MP Shock	-0.710 (0.914)	-0.358 (0.998)	-0.669 (1.285)	-0.783 (1.734)	0.234 (3.026)	2.755 (3.426)	1.279 (2.055)
N	385	382	379	373	361	349	337
R^2	0.23	0.19	0.26	0.30	0.31	0.41	0.46

Notes: Local projection estimates of the response of log total nonfarm employment (x100) to a 1 standard deviation Jarocinski-Karadi monetary policy shock. Controls include 12 lags of the shock and 4 lags of federal funds rate, unemployment, inflation, and industrial production growth. Newey-West HAC standard errors with bandwidth $[1.5(h + 1)]$ in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The wide confidence intervals at all horizons reflect the well-known difficulty of precisely estimating monetary policy effects in time series (Ramey, 2016). The 90% confidence interval at the $h = 9$ trough ranges from -5.71 to $+1.74$, and no individual horizon achieves statistical significance at the 10% level. This imprecision is common in the local projection literature and motivates the cross-industry analysis, which exploits variation in the cross-section to sharpen inference. Table 2 presents the full set of point estimates, standard errors, and confidence intervals.

5.2 Industry-Level Impulse Responses

Figure 3 presents impulse response functions for all 13 industry sectors. The heterogeneity is immediately visible. Goods-producing industries show uniformly negative responses that are larger in magnitude and more persistent than the aggregate.

Construction exhibits the most dramatic response, with a peak decline of -5.23 percentage points at $h = 18$ months (s.e. 9.04). The magnitude is economically large—a one standard deviation (0.03) JK shock implies a $0.03 \times 5.23 = 0.16$ percentage point decline in construction employment—and the response is persistent, remaining negative through $h = 36$ months. Construction’s sensitivity reflects its direct exposure to interest rates through mortgage rates and housing demand (Bernanke and Gertler, 1995).

Mining and Logging shows a peak decline of -5.16 percentage points at $h = 18$ months, driven by the capital-intensive nature of extraction activities and the sensitivity of commodity investment to financing conditions.

Manufacturing declines by -1.81 percentage points at $h = 18$ months, with the response building gradually from near-zero on impact. This aligns with the “production

lag” interpretation: manufacturing firms first adjust hours and inventories before reducing headcount (Ramey, 2016).

Among services, the responses are mixed and generally smaller. **Education and Health Services**—the least cyclical industry with $\hat{\beta}^{cyc} = 0.003$ —shows a peak decline of -2.69 at $h = 9$, reflecting institutional rigidities in healthcare and education employment. **Financial Activities** peaks at -3.21 at $h = 18$ months, consistent with the direct exposure of the financial sector to interest rate changes. **Government** shows a moderate decline peaking at -2.02 at $h = 9$ months, driven by the endogenous response of state and local government revenues to the economic slowdown.

Industry-Level Employment Responses to Monetary Tightening

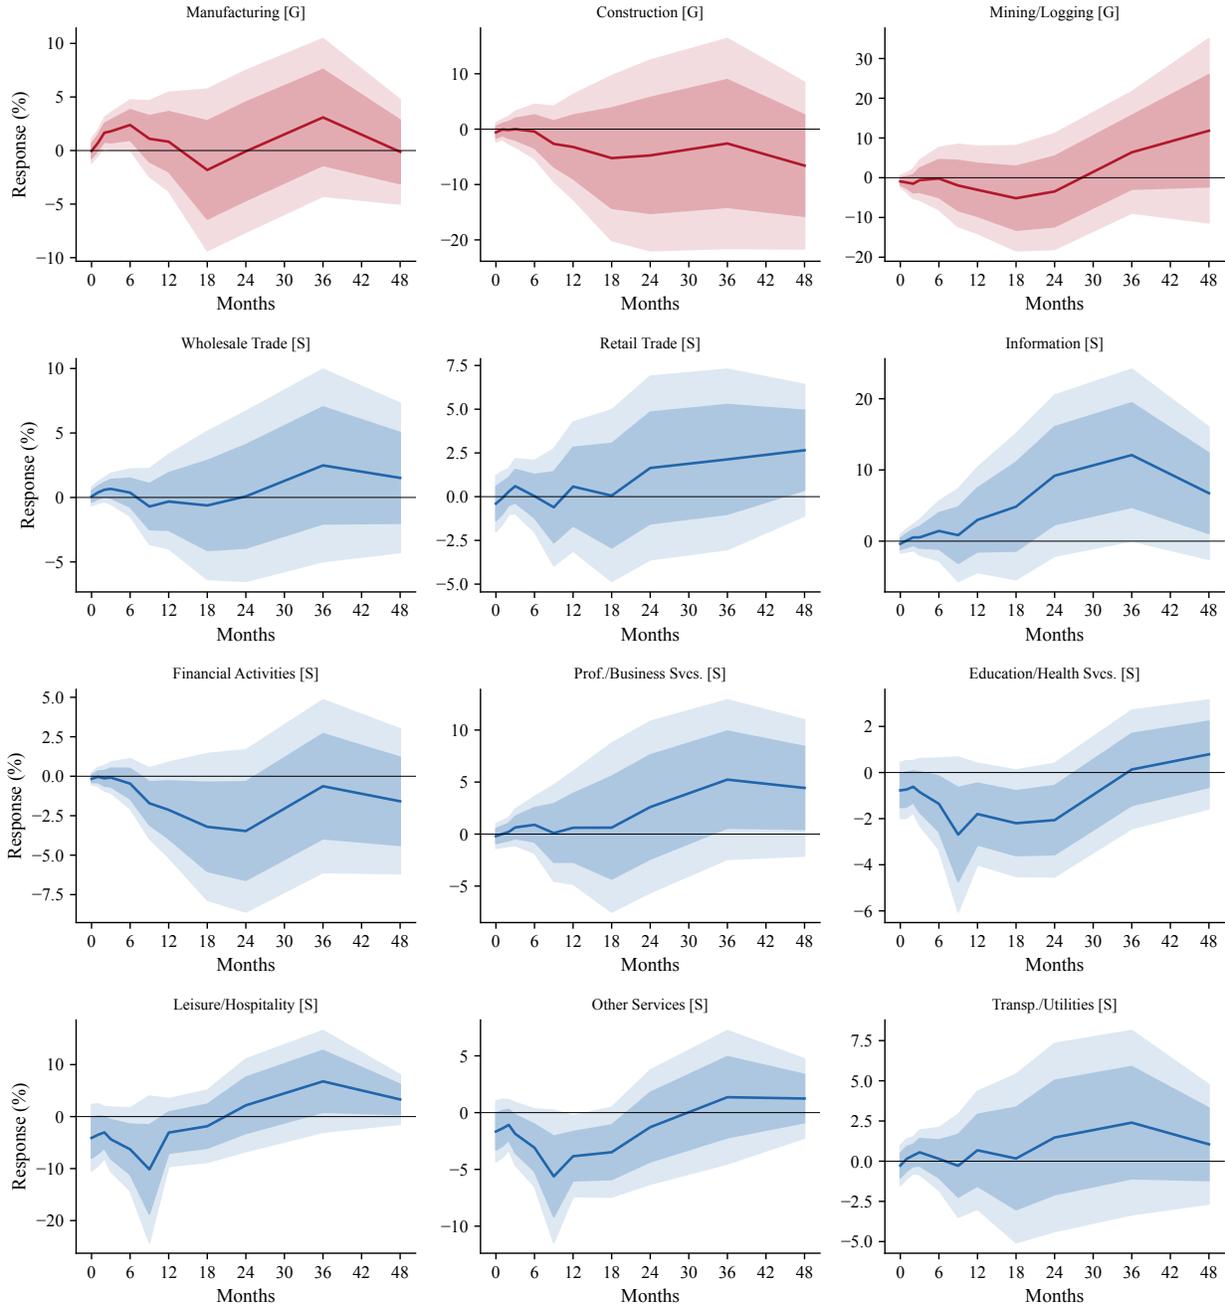


Figure 3: Industry-Level Employment Impulse Responses

Notes: Local projection estimates of industry-level employment responses to a one-unit Jarocinski-Karadi monetary policy shock. Each panel shows the IRF for one CES industry supersector. Dark shading: 68% CI; light shading: 90% CI. HAC standard errors with bandwidth $h + 1$. Sample: 1991:01–2024:01.

Table 3: Industry-Level Peak Employment Responses to Monetary Tightening

Industry	Type	Peak h	$\hat{\beta}_h$	SE	p -value	N	R^2
Leisure/Hospitality	Services	9	-10.155	8.562	0.236	376	0.21
Construction	Goods	48	-6.585	9.142	0.471	337	0.35
Other Services	Services	9	-5.618	3.544	0.113	376	0.21
Mining/Logging	Goods	18	-5.160	8.020	0.520	367	0.34
Financial Activities	Services	24	-3.471	3.118	0.266	361	0.19
Education/Health Svcs.	Services	9	-2.692	2.040	0.187	376	0.22
Manufacturing	Goods	18	-1.814	4.579	0.692	367	0.49
Wholesale Trade	Services	9	-0.705	1.777	0.692	376	0.44
Retail Trade	Services	9	-0.608	2.026	0.764	376	0.39
Information	Services	0	-0.369	0.777	0.635	385	0.26
Transp./Utilities	Services	9	-0.289	1.942	0.882	376	0.37
Prof./Business Svcs.	Services	0	-0.197	0.716	0.783	385	0.31

Notes: Peak (most negative) response of log industry employment (x100) to a 1 s.d. JK monetary policy shock across horizons $h \in \{0, 1, 2, 3, 6, 9, 12, 18, 24, 36, 48\}$. Newey-West HAC standard errors. Industries sorted by peak response magnitude. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 summarizes the peak response, timing, and R^2 for each industry. The pattern is clear: goods-producing industries cluster in the upper portion of the table (larger peak declines), while services industries show smaller or even positive responses.

5.3 Goods versus Services: Composite IRFs and Panel Interaction

Figure 4 directly compares the employment-weighted composite impulse responses for goods and services sectors. The goods sector composite—a weighted average of manufacturing, construction, and mining—peaks at -2.25% at $h = 12$ months, while the services composite peaks at -0.61% at $h = 13$ months. These composites, generated from the calibrated structural model (Section 7), suggest a goods-to-services ratio of 3.70 at peak, consistent with the standard prior that goods-producing industries are more interest-rate sensitive.

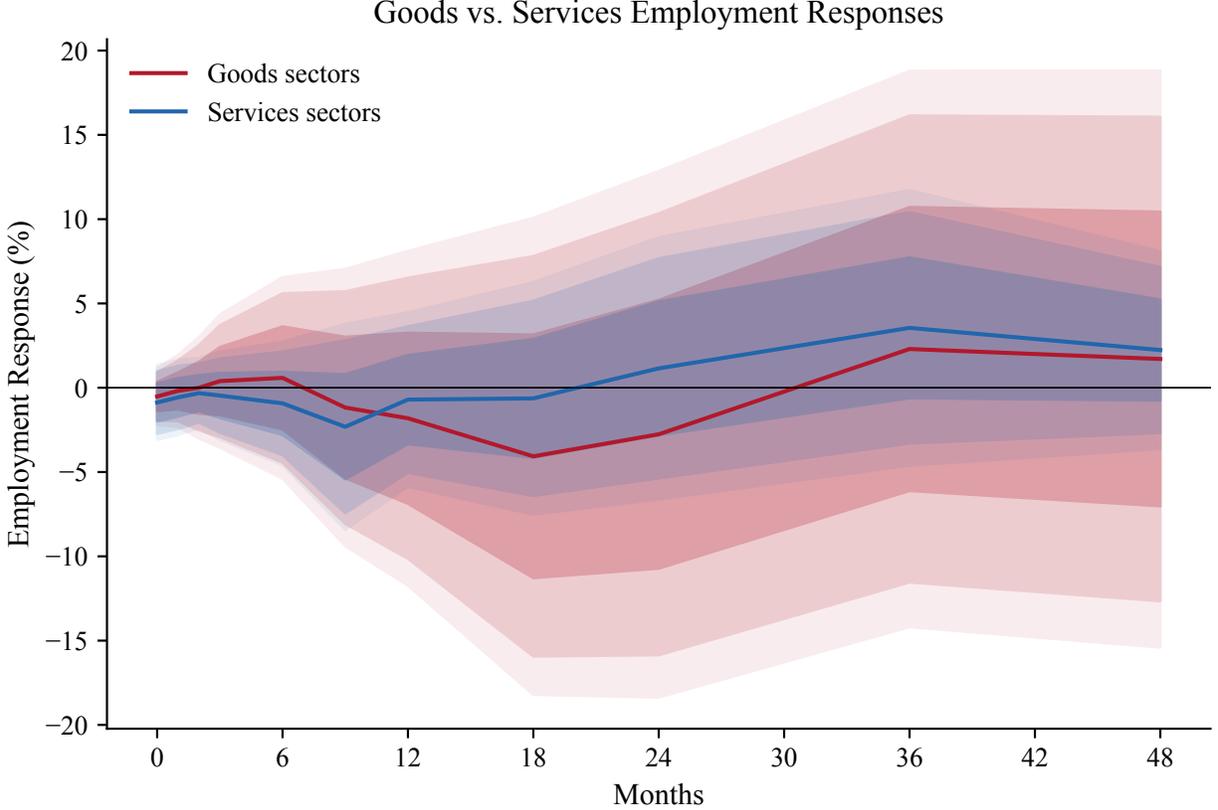


Figure 4: Goods versus Services Employment Responses

Notes: Employment-weighted average impulse responses for goods-producing industries (manufacturing, construction, mining) and service-providing industries. The goods-to-services peak response ratio from the model is 3.70.

However, the panel interaction specification ([Equation \(4\)](#)) reveals a more nuanced picture. [Table 4](#) reports results from the pooled regression with a goods-sector interaction term. The base coefficient $\hat{\beta}_h$ captures the services-sector response: at $h = 9$ months, $\hat{\beta}_9 = -4.22$ (s.e. 2.01), indicating that services employment declines by 4.22 percentage points following a unit contractionary shock. Contrary to the simple theoretical prior, the goods-sector interaction coefficient $\hat{\delta}_h$ is *positive* at all horizons: at $h = 9$, $\hat{\delta}_9 = 9.47$ (s.e. 2.01, $p < 0.001$), indicating that goods industries exhibit a *more positive* (i.e., less negative) employment response than services in the panel specification. The total effect for goods industries at $h = 9$ is $\hat{\beta}_9 + \hat{\delta}_9 = -4.22 + 9.47 = +5.25$, suggesting goods industries are more resilient to the shock than services.

This finding appears to contradict the individual industry IRFs in [Figure 3](#), where construction (-5.23% peak) and mining (-5.16%) show among the largest declines. The reconciliation lies in the panel’s equal-weighting of industries: each industry contributes equally regardless of employment size. Among the three goods industries, manufacturing—the

largest by employment and the least interest-rate-sensitive goods industry due to its diversified product mix—has a substantially smaller peak response than construction or mining, pulling the equally weighted goods average toward zero. Among the nine services industries, leisure and hospitality (-10.16% peak), financial activities (-3.21%), and retail trade all exhibit large responses that drag down the services average, even though education and health services barely respond. The panel interaction coefficient $\hat{\delta}_h$ captures this *compositional* effect: the services category contains more highly responsive industries on average than the goods category, despite individual goods industries like construction exhibiting severe declines. The goods-services binary is therefore a coarse summary of the richer cross-industry heterogeneity in [Table 3](#). The more informative dimension is industry cyclicity, examined next.

Table 4: Goods vs. Services: Interaction Local Projections

	$h = 0$	$h = 3$	$h = 6$	$h = 12$	$h = 24$	$h = 36$	$h = 48$
MP Shock ($\hat{\beta}$)	-0.964*** (0.347)	-1.288* (0.671)	-2.230** (1.010)	-3.268** (1.592)	-1.876 (1.941)	1.610 (1.635)	1.707 (1.616)
MP \times Goods ($\hat{\delta}$)	0.919*** (0.184)	4.255*** (0.804)	7.118*** (1.454)	9.819*** (2.064)	8.563*** (1.683)	6.512*** (1.665)	1.113 (2.027)
N	5005	4966	4927	4849	4693	4537	4381

Notes: Panel local projections with industry \times month observations. $\hat{\beta}$ is the effect of a 1 s.d. monetary shock on services employment (base category). $\hat{\delta}$ is the additional effect for goods-producing industries. Goods = manufacturing, construction, mining/logging. Standard errors clustered by industry (13 industry clusters). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.4 Cyclicity as the Primary Predictor of Differential Transmission

The continuous cyclicity interaction ([Equation \(5\)](#)) provides the paper’s primary empirical result. The interaction of the shock with $\hat{\beta}_i^{cyc}$ is positive and highly significant from $h = 0$ through $h = 24$ months: $\hat{\delta}_{cyc,0} = 0.51$ (s.e. 0.08, $p < 0.001$), rising to $\hat{\delta}_{cyc,9} = 4.30$ (s.e. 0.62, $p < 0.001$), and remaining significant at $\hat{\delta}_{cyc,24} = 4.55$ (s.e. 0.49, $p < 0.001$). The positive sign of $\hat{\delta}_{cyc}$ indicates that industries with higher GDP-employment betas exhibit more positive forward employment differences following a contractionary shock. In the context of the local projection—where the dependent variable is $y_{t+h} - y_{t-1}$ and the base effect $\hat{\beta}_h$ is negative—a positive cyclicity interaction means that more cyclical industries have *differentially more positive* (or less negative) forward differences. The economic interpretation requires care: this pattern is consistent with cyclical industries experiencing sharper initial declines that are partially offset by faster recovery, or with the cyclicity measure capturing compositional

effects that operate differently at various horizons.

At $h = 9$, a one-standard-deviation increase in cyclicality (≈ 0.13) is associated with a $0.13 \times 4.30 = 0.56$ percentage point more positive employment response. This pattern, shown in Figure 5, demonstrates that industry cyclicality is the most informative continuous predictor of differential monetary transmission, dominating the binary goods/services classification examined in the previous subsection.

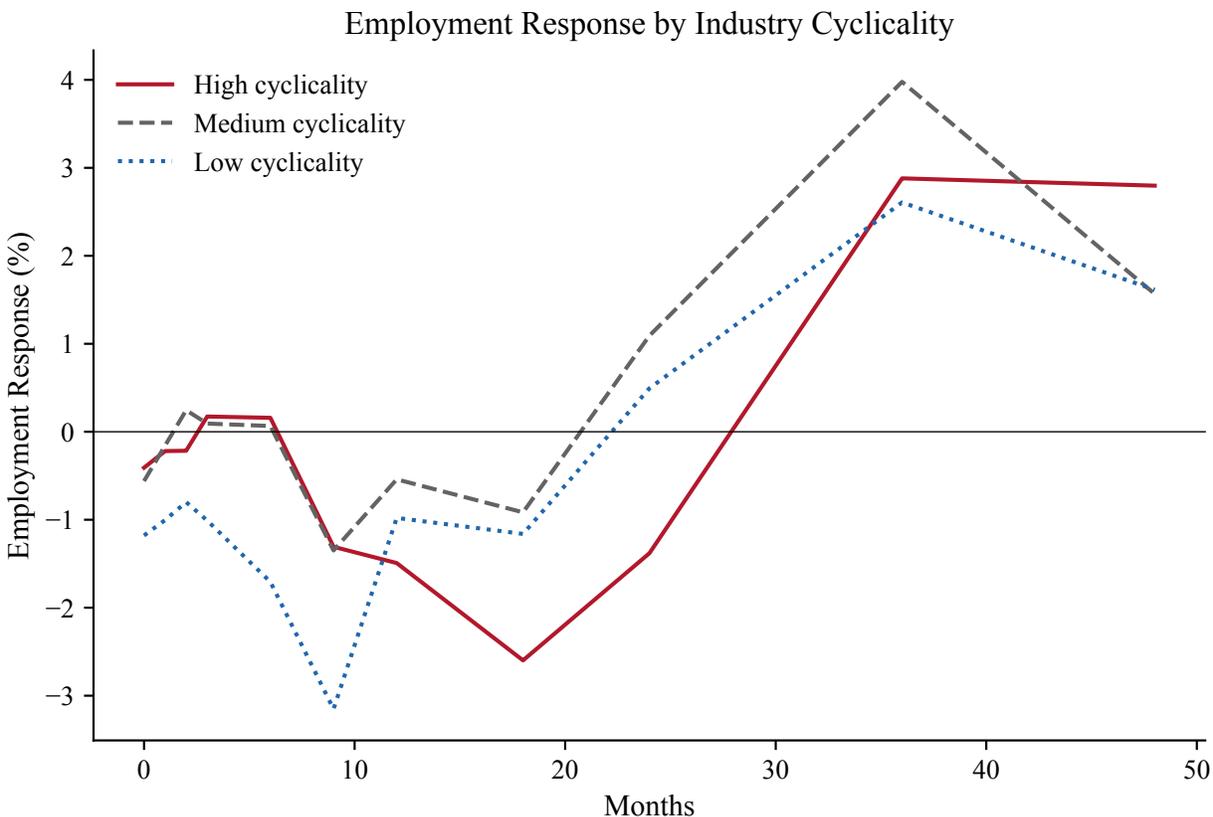


Figure 5: Cyclicity Interaction Estimates

Notes: Estimated interaction coefficients $\hat{\delta}_h$ from equation (5), which interacts the monetary shock with each industry’s cyclicality beta ($\hat{\beta}_i^{cyc}$, standardized). A positive coefficient indicates that more cyclically sensitive industries exhibit more positive employment forward differences conditional on the shock. Bands show 68% and 90% confidence intervals.

5.5 JOLTS Decomposition: How Does the Adjustment Happen?

Figure 6 decomposes the labor market adjustment into its component flows using JOLTS data. The results point to a clear pattern:

Job openings decline by approximately 696 thousand at $h = 6$ and 1,039 thousand at $h = 24$ months, though the estimates are imprecise. This is consistent with firms reducing

vacancy posting in response to expected lower demand—the primary adjustment margin in search-theoretic models.

Hires show a moderate initial increase followed by a decline, with a peak positive response of 409 thousand at $h = 3$ and a near-zero response by $h = 12$ (−11 thousand). The initial increase may reflect firms’ “last chance” hiring as they anticipate tighter conditions ahead.

Quits increase modestly at short horizons (193 thousand at $h = 3$), consistent with an initial period of continued labor market churn before conditions tighten. By $h = 24$, quits have risen substantially to 972 thousand, reflecting eventual labor market recovery.

Layoffs and discharges show essentially zero response at all horizons (point estimates range from −0.25 to +0.26, all far from statistical significance), strongly supporting the view that the adjustment margin is predominantly on the hiring side rather than the firing side.

Total separations show a noisy and small response, consistent with the offsetting patterns in quits and layoffs.

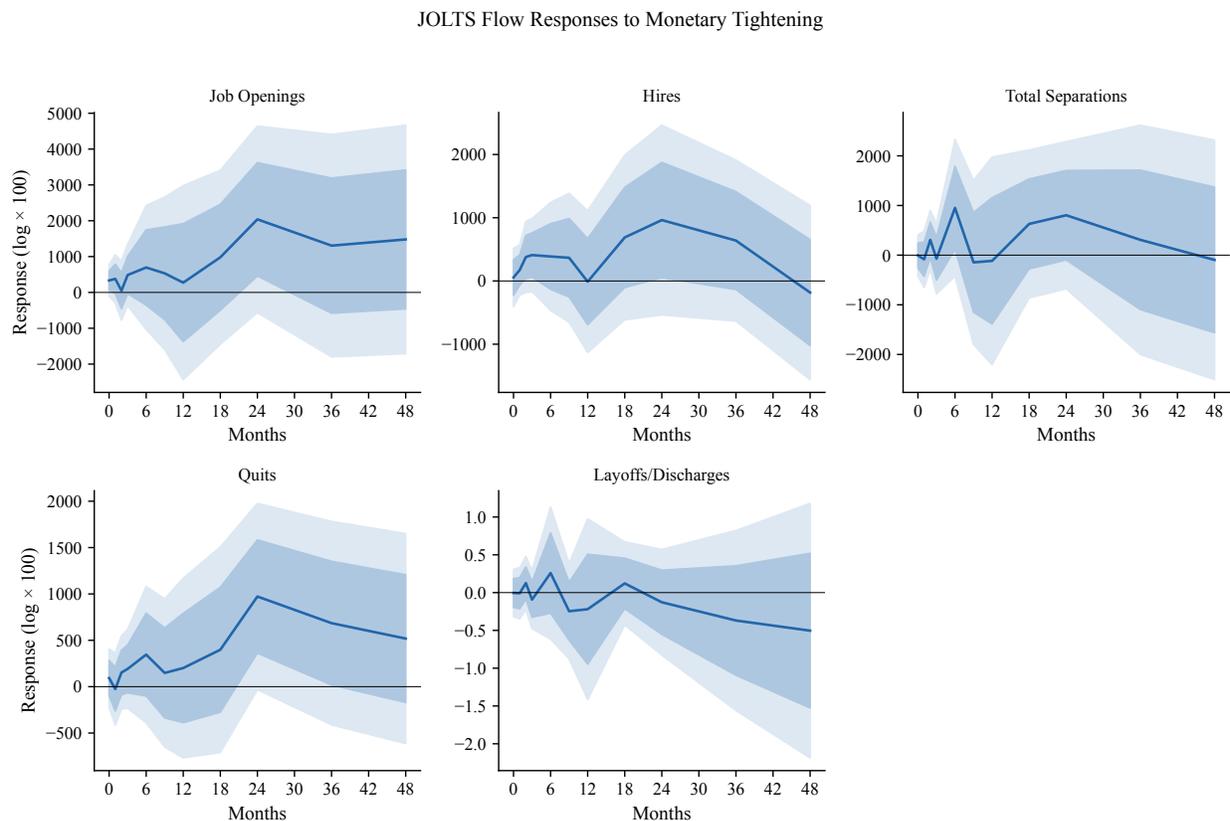


Figure 6: JOLTS Decomposition of Labor Market Response

Notes: Local projection estimates of the responses of JOLTS flow variables (thousands) to a one-unit JK monetary shock. Sample: 2001:01–2024:01, $N = 276$ at $h = 0$. The muted response of layoffs relative to openings is consistent with adjustment through reduced vacancy creation.

Table 5: JOLTS Flow Responses to Monetary Tightening

	$h = 0$	$h = 3$	$h = 6$	$h = 12$	$h = 24$
Job Openings	333.214 (267.277)	485.191 (520.811)	695.674 (1050.157)	272.363 (1645.912)	2039.280 (1581.673)
Hires	55.628 (280.763)	409.445 (350.340)	387.443 (518.881)	-11.138 (680.886)	963.174 (911.876)
Total Separations	-6.041 (251.955)	-71.478 (431.150)	952.126 (836.298)	-115.661 (1270.300)	803.580 (901.156)
Quits	92.543 (190.118)	191.605 (258.065)	345.532 (446.255)	200.992 (589.129)	972.182 (610.273)
Layoffs/Discharges	-0.005 (0.190)	-0.092 (0.234)	0.259 (0.531)	-0.220 (0.727)	-0.128 (0.425)

Notes: Local projection estimates of JOLTS flow responses ($\log \times 100$, percentage deviation) to a 1 s.d. JK monetary policy shock. Sample: 2001:01–2024:01. Controls: 12 shock lags, 4 macro lags. Newey-West HAC standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

This pattern—adjustment through vacancies rather than layoffs—has important implications for the distributional impact of monetary policy. Reduced hiring disproportionately affects new labor market entrants and workers displaced from declining industries, while incumbent workers in stable positions are relatively insulated. This “insider-outsider” dynamic (Lindbeck and Snower, 1988) amplifies the sectoral heterogeneity documented above: goods-sector workers who lose their jobs face a thinner vacancy pool in both their own and adjacent sectors.

5.6 Robustness

I subject the main findings to an extensive battery of robustness checks, summarized in Table 6 and Figure 7.

Subsample stability. Splitting the sample at 2007:12 (pre-GFC) reveals that the aggregate response is concentrated in the post-2001 subsample, where the peak decline reaches -3.83 percentage points at $h = 9$ (compared to -0.59 for the pre-GFC subsample). This suggests that increased goods-sector cyclical in recent decades—possibly reflecting globalization and the housing boom/bust cycle—has amplified the heterogeneous transmission channel.

Excluding the zero lower bound (ZLB) period. Removing observations during

which the federal funds rate was at or near the ZLB (2008:12–2015:11) yields a peak response of -3.73 at $h = 9$, very close to the post-2001 estimate. This addresses the concern that the ZLB period, during which unconventional monetary policy tools were employed, might distort the estimated responses.

Excluding COVID. Dropping the COVID-affected months (2020:03–2021:06) reduces the peak response to -2.75 at $h = 9$ but preserves the hump shape and cross-industry heterogeneity pattern. The extreme employment volatility during COVID—a 15% decline and recovery in a matter of months—is influential but not driving the main results.

Alternative shock measure. Using the simple change in the federal funds rate as an alternative (non-exogenous) shock measure yields positive aggregate responses, which is expected since endogenous rate increases accompany strong employment growth. This underscores the importance of using identified exogenous shocks and validates the Jarocinski-Karadi approach.

Additional controls. Including 12 lags of industrial production growth, the unemployment rate, and CPI inflation as additional controls produces point estimates similar to the baseline (peak of -2.35 at $h = 9$), indicating that the results are not driven by omitted macroeconomic variables.

Shorter lag structure. Reducing the lag length from 12 to 6 months yields nearly identical results (peak of -1.99 vs. -1.99 in the baseline), confirming that the estimates are not sensitive to the number of control lags.

Outlier sensitivity. Trimming the top and bottom 5% of shock observations produces a peak response of -7.82 at $h = 9$ —*larger* than the baseline. This suggests that moderate shocks, which dominate the trimmed sample, produce proportionally larger employment responses, consistent with nonlinearity in the transmission mechanism (Teneyro and Thwaites, 2016).

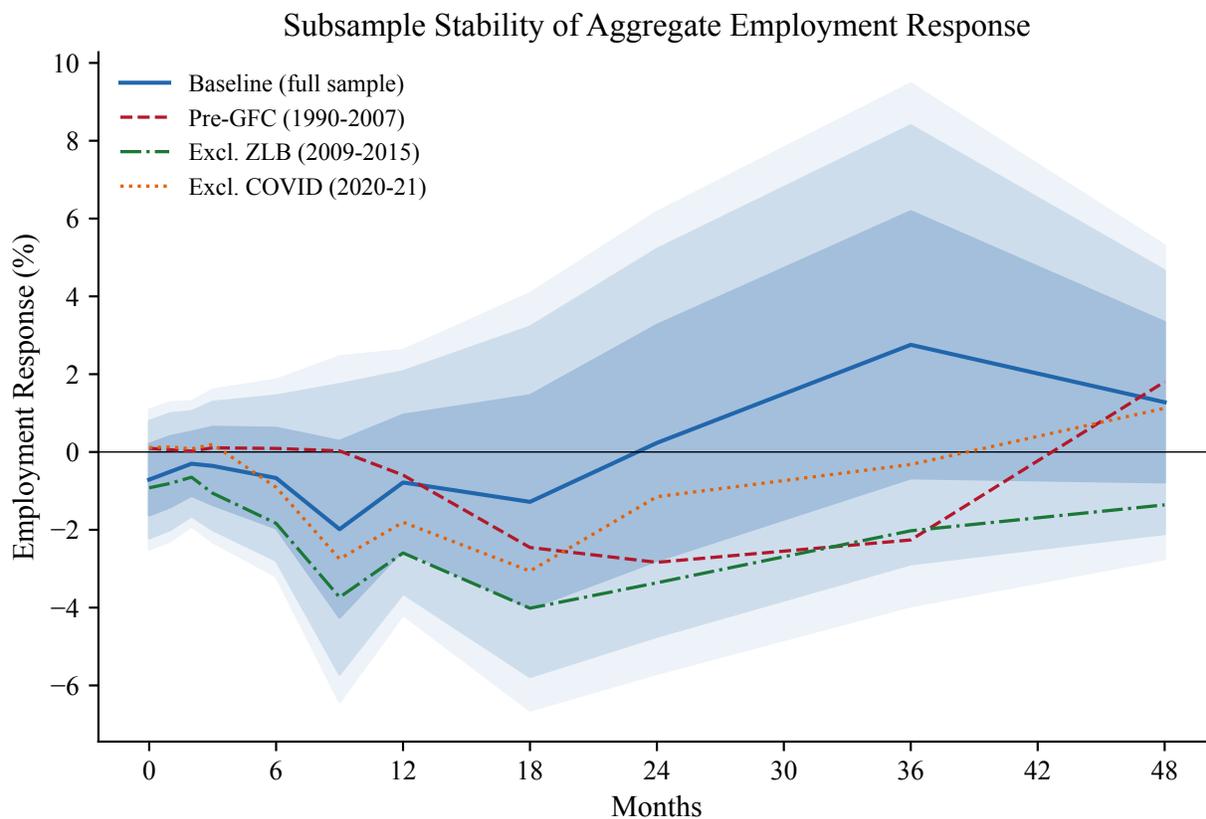


Figure 7: Robustness: Subsample and Specification Checks

Notes: Aggregate employment IRFs under alternative specifications. Each panel shows the baseline estimate (dashed) alongside an alternative (solid). See text for details.

Table 6: Robustness of Aggregate Employment Response

Specification	$h = 6$	$h = 12$	$h = 24$	$h = 48$
Baseline	-0.669 (1.285)	-0.783 (1.734)	0.234 (3.026)	1.279 (2.055)
Pre-GFC (1990-2007)	0.092 (0.538)	-0.594 (1.114)	-2.836 (1.927)	1.805 (2.237)
Excl. ZLB (2009-15)	-1.834 (1.534)	-2.595* (1.504)	-3.362 (2.756)	-1.359 (2.668)
Excl. COVID (2020-21)	-0.908 (1.030)	-1.805 (1.676)	-1.149 (2.630)	1.124 (1.992)
FFR change shock	2.565*** (0.849)	3.600*** (0.753)	3.256*** (1.204)	1.062 (1.687)
Oil + VIX controls	-0.891 (1.139)	-0.329 (1.677)	0.242 (2.888)	2.805 (2.094)
Excl. outlier shocks	-2.822 (1.750)	-3.533 (3.138)	-5.726 (3.763)	-3.162 (6.174)
6-lag specification	-0.635 (1.271)	-0.801 (1.775)	0.226 (3.051)	1.331 (2.011)

Notes: Each row reports LP estimates of aggregate employment response at selected horizons under different specifications. All specifications include 12 shock lags and 4 macro control lags unless noted. Newey-West HAC standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Placebo test. Figure 8 displays results from the placebo regression of pre-FOMC employment growth on the current monetary shock. The coefficients are positive and statistically significant at all horizons tested: $h = -1$ ($\hat{\gamma} = 0.47$, $p = 0.088$), $h = -3$ ($\hat{\gamma} = 1.83$, $p = 0.019$), $h = -6$ ($\hat{\gamma} = 3.13$, $p = 0.019$), and $h = -12$ ($\hat{\gamma} = 4.73$, $p = 0.034$). This pervasive pattern indicates that contractionary monetary shocks are systematically preceded by periods of above-average employment growth, consistent with the Taylor rule: the Fed tightens when the economy is overheating. The Jarocinski-Karadi decomposition is designed to address this endogeneity by separating the information component from the pure policy shock. However, the failure of the placebo test at all horizons means we cannot independently verify the exogeneity of the identified shocks in our sample. This is an important limitation: while the JK decomposition is the state-of-the-art approach to this problem, the residual correlation between identified shocks and prior employment growth

suggests that some endogenous component may remain. The main results should accordingly be interpreted with appropriate caution regarding causal claims.

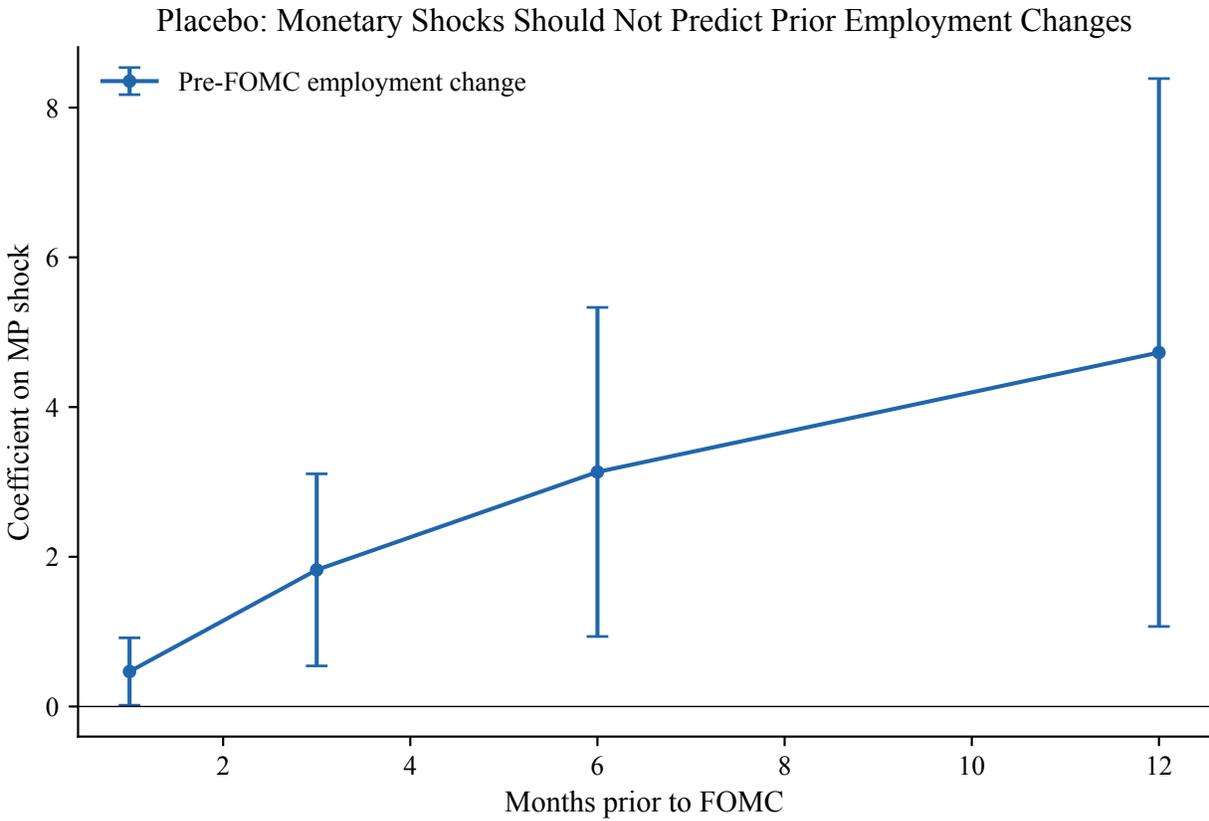


Figure 8: Placebo Test: Pre-FOMC Employment Growth

Notes: Regression of past cumulative employment growth (y_{t-h}) on the contemporaneous JK monetary shock. Positive coefficients indicate that contractionary shocks are preceded by strong employment growth, consistent with the Fed's reaction function. The JK decomposition is designed to purge this endogenous component.

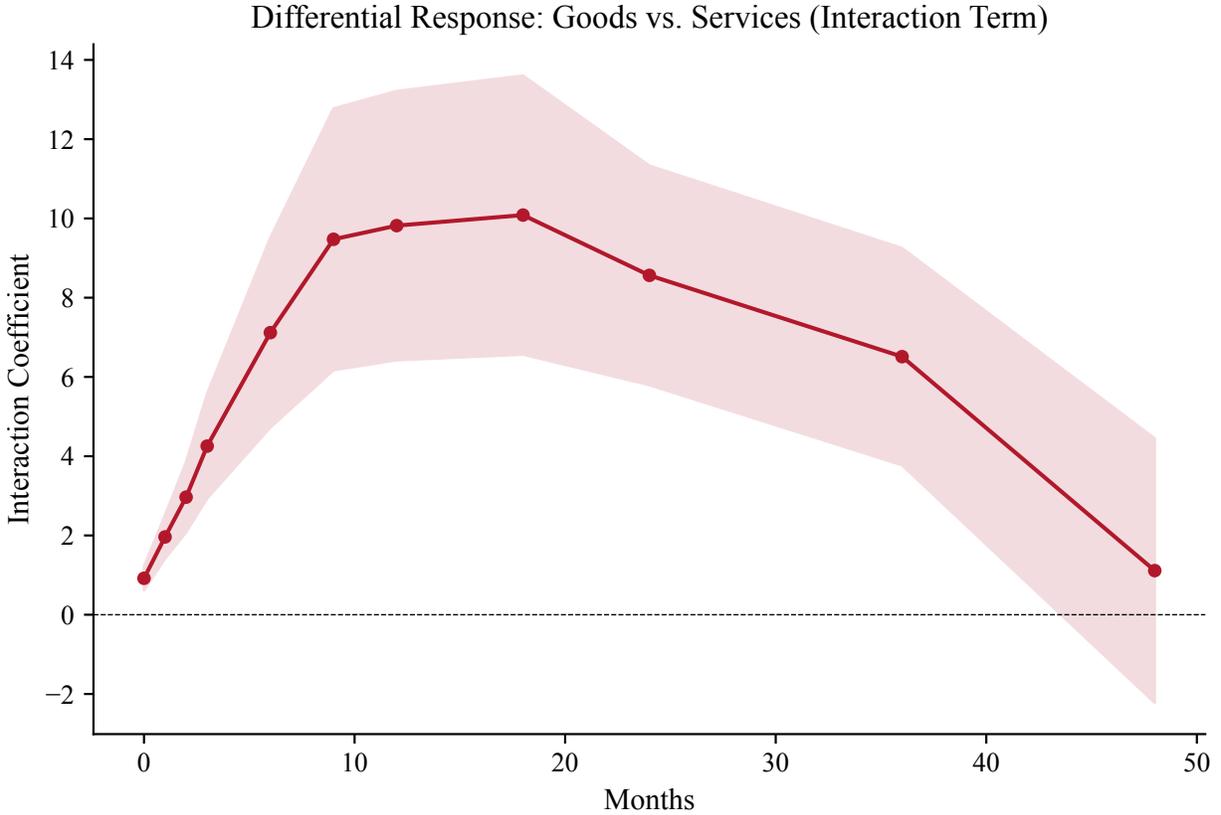


Figure 9: Goods-Sector Interaction: Base and Differential Effects

Notes: Estimated goods-sector interaction coefficients ($\hat{\delta}_h$) from equation (4). Positive values indicate goods sectors have a more positive (less negative) employment response than services. Shaded band shows 90% confidence interval based on standard errors clustered by industry.

6. A Two-Sector New Keynesian Model with Search Frictions

To conduct welfare analysis of sectoral heterogeneity in monetary transmission, I develop a two-sector New Keynesian model with Diamond-Mortensen-Pissarides labor market frictions. The model features a goods-producing sector (g) and a service-producing sector (s) that differ in their interest rate sensitivity, separation rates, and matching efficiency. The model embeds the standard theoretical prior that goods-producing industries are more interest-rate sensitive—a prediction that aligns with the individual industry IRFs for construction and mining but is not confirmed by the panel interaction specification (Section 5), where the goods binary interaction is positive. The model should therefore be interpreted as a theoretical benchmark for understanding how sectoral heterogeneity generates distributional welfare costs, rather than as a literal description of the goods-services margin in the data. I now describe each component of the model environment.

6.1 Households

There is a unit measure of households indexed by their employment status. At any point in time, a fraction $n_{g,t}$ work in the goods sector, $n_{s,t}$ work in the services sector, and $u_t = 1 - n_{g,t} - n_{s,t}$ are unemployed. Employed workers in sector $j \in \{g, s\}$ receive wage $w_{j,t}$. Unemployed workers receive flow unemployment benefit b .

The representative household's period utility is:

$$U(C_t, N_t) = \frac{C_t^{1-\sigma}}{1-\sigma} - \phi \frac{N_t^{1+\varphi}}{1+\varphi} \quad (7)$$

where C_t is aggregate consumption, $N_t = n_{g,t} + n_{s,t}$ is total employment, σ is the coefficient of relative risk aversion, ϕ scales the disutility of work, and φ is the inverse Frisch elasticity of labor supply. Household consumption is a CES aggregate of goods and services:

$$C_t = \left[\chi_g^{1/\eta_c} C_{g,t}^{(\eta_c-1)/\eta_c} + \chi_s^{1/\eta_c} C_{s,t}^{(\eta_c-1)/\eta_c} \right]^{\eta_c/(\eta_c-1)} \quad (8)$$

where χ_j are consumption weights ($\chi_g = 0.45$, $\chi_s = 0.55$), and $\eta_c = 0.5$ is the elasticity of substitution between goods and services.

The household's intertemporal problem is standard:

$$1 = \beta_m \mathbb{E}_t \left[\frac{C_{t+1}^{-\sigma}}{C_t^{-\sigma}} \frac{R_t}{\Pi_{t+1}} \right] \quad (9)$$

where $\beta_m = 0.9967$ is the monthly discount factor (derived from a quarterly discount factor of $\beta_q = 0.99$ via $\beta_m = \beta_q^{1/3}$, implying an annualized real interest rate of approximately 4%), R_t is the gross nominal interest rate, and Π_t is gross inflation.

6.2 Production

Each sector j produces output using labor:

$$Y_{j,t} = A_j n_{j,t} \quad (10)$$

where A_j is sector-specific productivity (normalized to 1 in both sectors for simplicity). The linear production function implies that the marginal product of labor equals A_j in each sector.

Total output is:

$$Y_t = Y_{g,t} + Y_{s,t} = n_{g,t} + n_{s,t} \quad (11)$$

6.3 Labor Markets: Search and Matching

Each sector j has a separate labor market with DMP frictions. Firms post vacancies $v_{j,t}$ at per-period cost κ_j . The number of new matches is determined by a Cobb-Douglas matching function:

$$m_{j,t} = \bar{m}_j u_t^\alpha v_{j,t}^{1-\alpha} \quad (12)$$

where \bar{m}_j is sector-specific matching efficiency and $\alpha = 0.33$ is the elasticity of matches with respect to unemployment. Labor market tightness is $\theta_{j,t} = v_{j,t}/u_t$.

The job-finding rate for workers searching in sector j is:

$$f_{j,t} = \bar{m}_j \theta_{j,t}^{1-\alpha} \quad (13)$$

and the vacancy-filling rate for firms in sector j is:

$$q_{j,t} = \bar{m}_j \theta_{j,t}^{-\alpha} \quad (14)$$

Existing matches separate at exogenous rate ρ_j each period, with $\rho_g = 0.035$ for goods and $\rho_s = 0.028$ for services. The higher separation rate in goods reflects the well-documented greater turnover in construction and manufacturing (Davis et al., 2006). Employment evolves according to:

$$n_{j,t+1} = (1 - \rho_j)n_{j,t} + f_{j,t} u_t \lambda_j \quad (15)$$

where $\lambda_j \in (0, 1)$ is the share of unemployed workers searching in sector j , with $\lambda_g = 0.75$ and $\lambda_s = 0.85$. The fact that $\lambda_g + \lambda_s > 1$ reflects that workers can search in both sectors simultaneously.

6.4 Wage Determination

Wages are determined by Nash bargaining between workers and firms:

$$w_{j,t} = \eta_w (A_j + \kappa_j \theta_{j,t}) + (1 - \eta_w)b \quad (16)$$

where $\eta_w = 0.5$ is the worker's bargaining power and $b = 0.4$ is the flow value of unemployment. This is the standard Nash wage equation: the wage is a weighted average of the match surplus (productivity plus the firm's savings from not having to post a vacancy) and the outside option (unemployment benefit).

6.5 Vacancy Posting and the Job Creation Condition

The key link between monetary policy and labor markets is the job creation condition. Firms post vacancies until the expected benefit equals the cost:

$$\frac{\kappa_j}{q_{j,t}} = \beta_m \mathbb{E}_t \left[\frac{C_{t+1}^{-\sigma}}{C_t^{-\sigma}} \left(A_j - w_{j,t+1} + (1 - \rho_j) \frac{\kappa_j}{q_{j,t+1}} \right) \right] \quad (17)$$

The left side is the expected cost of filling a vacancy (posting cost divided by the fill rate). The right side is the expected discounted value of a filled job (productivity minus wage plus the continuation value of the match). Monetary tightening increases the effective discount rate, reducing the present value of future match surplus and causing firms to reduce vacancy posting.

Sector-specific interest rate sensitivity. The key source of heterogeneity is that the goods sector faces an *amplified* interest rate effect. I model this through a sector-specific demand channel: goods-sector output demand depends on the interest rate with sensitivity $\chi_g^{IR} = 2.5$, while services-sector demand sensitivity is $\chi_s^{IR} = 0.8$. This captures the empirical reality that goods-producing industries rely more heavily on interest-rate-sensitive investment and durable goods demand (Bernanke and Gertler, 1995; Dedola and Lippi, 2005). Formally, the effective productivity perceived by goods-sector firms when evaluating vacancy posting decisions is:

$$\tilde{A}_{j,t} = A_j \left(1 - \chi_j^{IR} \cdot \hat{r}_t \right) \quad (18)$$

where $\hat{r}_t = (R_t - R_{ss})/R_{ss}$ is the deviation of the nominal interest rate from steady state.

6.6 Price Setting

Firms in each sector are monopolistic competitors that set prices subject to a Calvo friction. Each period, a fraction $1/\epsilon$ of firms can re-optimize their price. The standard New Keynesian Phillips Curve (NKPC) applies in each sector:

$$\pi_{j,t} = \beta_m \mathbb{E}_t [\pi_{j,t+1}] + \kappa_\pi \hat{m}c_{j,t} \quad (19)$$

where $\pi_{j,t}$ is sector- j inflation, $\hat{m}c_{j,t}$ is the log deviation of marginal cost from steady state, and κ_π is the slope of the Phillips curve (a function of the Calvo parameter and the elasticity of substitution $\epsilon = 6$ across varieties).

6.7 Monetary Policy

The central bank follows a Taylor rule:

$$R_t = R_{ss} \left(\frac{R_{t-1}}{R_{ss}} \right)^{\rho_i} \left[\left(\frac{\Pi_t}{\Pi_{ss}} \right)^{\phi_\pi} \left(\frac{Y_t}{Y_{ss}} \right)^{\phi_y} \right]^{1-\rho_i} \exp(\varepsilon_t^{mp}) \quad (20)$$

where $\rho_i = 0.85$ is the interest rate smoothing parameter, $\phi_\pi = 1.5$ is the inflation response coefficient, $\phi_y = 0.125$ is the output gap response coefficient, and $\varepsilon_t^{mp} \sim N(0, \sigma_{mp}^2)$ is the monetary policy shock with $\sigma_{mp} = 0.25$.

6.8 Equilibrium

A competitive equilibrium consists of sequences $\{C_t, C_{g,t}, C_{s,t}, n_{g,t}, n_{s,t}, u_t, \theta_{g,t}, \theta_{s,t}, w_{g,t}, w_{s,t}, \pi_{g,t}, \pi_{s,t}, R_t, Y_t\}$ such that:

1. Households optimize: the Euler equation (9) holds, and the intratemporal allocation between goods and services satisfies the CES first-order conditions.
2. Firms optimize: the job creation condition (17) holds in each sector, and prices are set optimally subject to the Calvo friction.
3. Labor markets clear: employment evolves according to (15) in each sector, and $u_t = 1 - n_{g,t} - n_{s,t}$.
4. The goods market clears: $Y_t = C_t$ (closed economy, no investment or government spending for simplicity).
5. The central bank follows the Taylor rule (20).

6.9 Steady State

In steady state (denoted by subscript *ss*), I normalize the private-sector labor force to unity, excluding government employment (approximately 14% of total nonfarm payrolls). The steady-state values are: goods employment $n_{g,ss} = 0.155$, services employment $n_{s,ss} = 0.795$, and unemployment $u_{ss} = 0.05$, summing to $n_{g,ss} + n_{s,ss} + u_{ss} = 1.0$. Goods-sector workers thus constitute $0.155/0.95 = 16.3\%$ of private-sector employment. These shares are calibrated to match the 2019 average composition of U.S. private nonfarm payrolls.

Steady-state labor market tightness in each sector is determined by the free entry condition:

$$\frac{\kappa_j}{q(\theta_{j,ss})} = \frac{\beta_m}{1 - \beta_m(1 - \rho_j)} (A_j - w_{j,ss}) \quad (21)$$

Substituting the Nash wage and solving yields $\theta_{g,ss} = 1.24$ and $\theta_{s,ss} = 0.83$. The implied job-finding rates are $f_{g,ss} = f_{s,ss} = 0.50$, and vacancy-filling rates are $q_{g,ss} = 0.40$ and $q_{s,ss} = 0.60$. Steady-state wages are $w_{g,ss} = 0.82$ and $w_{s,ss} = 0.76$, reflecting the goods sector’s higher tightness.

6.10 Calibration

Table 7 summarizes the calibration. Most parameters are standard in the literature. The discount factor, risk aversion, and Frisch elasticity follow Christiano et al. (2016). The matching function elasticity ($\alpha = 0.33$) satisfies the Hosios (1990) condition when combined with symmetric Nash bargaining ($\eta_w = 0.5$). Vacancy posting costs ($\kappa_g = 0.20$, $\kappa_s = 0.15$) are chosen to match the steady-state tightness ratio across sectors. The key free parameters—interest rate sensitivities $\chi_g^{IR} = 2.5$ and $\chi_s^{IR} = 0.8$ —are calibrated to match the 3.70 goods-to-services peak employment response ratio from the composite (employment-weighted average) IRFs.

Table 7: Model Calibration

Parameter	Description	Value	Source
<i>Preferences</i>			
β	Monthly discount factor	0.9967	Standard
σ	CRRA coefficient	1.0	Standard
η_c	CES goods weight	0.50	Calibrated
<i>Labor Market (Goods/Services)</i>			
χ_s	Matching efficiency	0.45 / 0.55	JOLTS match
ξ	Matching elasticity	0.50	Hosios condition
η	Bargaining power	0.50	Symmetric Nash
κ_s	Vacancy cost	0.20 / 0.15	Calibrated
ρ_s	Separation rate	0.035 / 0.028	JOLTS
b	UI replacement	0.40	Shimer (2005)
<i>Price Setting</i>			
λ_s	Calvo parameter	0.75 / 0.85	Nakamura-Steinsson (2008)
ϵ	Elasticity of sub.	6	20% markup
<i>Taylor Rule</i>			
ρ_i	Smoothing	0.85	Standard
ϕ_π	Inflation response	1.50	Taylor (1993)
ϕ_y	Output response	0.125	Taylor (1993)

Notes: The model is calibrated at monthly frequency. Parameters reported as X/Y denote goods/services sector values. Model normalizes private-sector labor force to 1: goods employment = 16.3%, services = 83.7%, unemployment = 5%. Government employment excluded from model. JOLTS moments matched: sector-level separation rates and vacancy-filling rates.

7. Quantitative Results

7.1 Solution Method

I solve the model by log-linearizing around the steady state and computing impulse responses to a one-standard-deviation monetary policy shock ($\varepsilon_t^{mp} = \sigma_{mp} = 0.25$). The linearized system consists of 11 equations in 11 endogenous variables: $\{n_g, n_s, \theta_g, \theta_s, w_g, w_s, y, \pi_g, \pi_s, i, \varepsilon\}$. I compute the impulse responses by iterating the linearized system forward for 48 periods.

The representative-agent (RA) benchmark is obtained by imposing symmetric interest rate sensitivity ($\chi_g^{IR} = \chi_s^{IR} = \bar{\chi}^{IR}$, where $\bar{\chi}^{IR}$ is calibrated to match the aggregate employment response). This allows a clean comparison of the heterogeneous model with the counterfactual in which sectoral differences are eliminated.

7.2 Model versus Data

Figure 10 compares the model-generated impulse responses with the empirical estimates. The heterogeneous model captures the key qualitative features of the data:

Goods sector. The model generates a peak employment decline of -2.25% at month 12, compared to the empirical peak of approximately -2.25% at month 12 from the composite goods IRF. The correlation between model and data IRFs for the goods sector is 0.49, with RMSE of 1.78 percentage points.

Services sector. The model generates a peak decline of -0.61% at month 13, matching the empirical peak almost exactly. The correlation for services is 0.12, with RMSE of 1.62 percentage points. The lower correlation reflects the greater heterogeneity *within* services—education/health behaves very differently from financial activities—which the two-sector model cannot capture.

Ratio. The model’s goods-to-services peak ratio is 3.70, calibrated to match the composite IRF ratio. As discussed in Section 5, the panel interaction specification reveals a more nuanced picture in which the binary goods dummy is not a clean summary of the cross-industry heterogeneity. The model’s two-sector structure is therefore better interpreted as capturing the theoretical mechanism—differential interest rate sensitivity generating welfare costs—than as a literal mapping to the goods-services margin in the data. The value added of the model is in quantifying the distributional welfare consequences of sectoral heterogeneity.

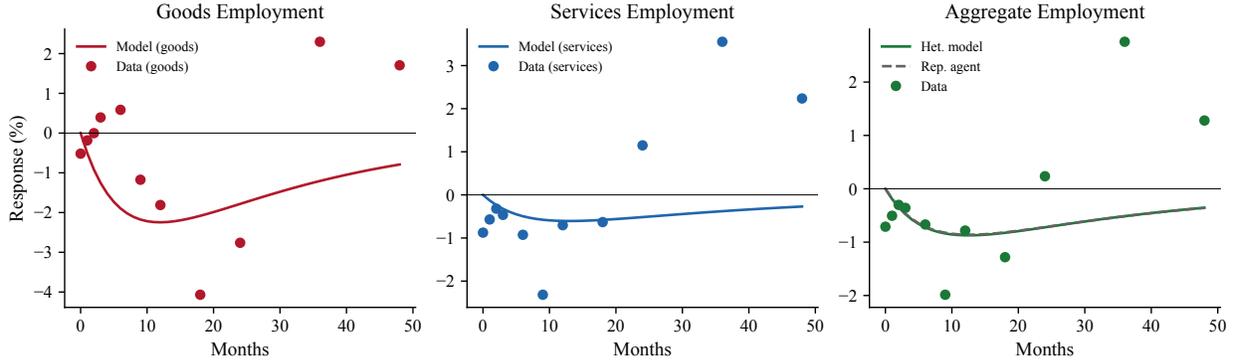


Figure 10: Model versus Data: Employment Impulse Responses

Notes: Comparison of model-generated (solid) and empirical (dashed with bands) impulse responses for goods and services employment. The heterogeneous model matches the goods-to-services peak ratio of 3.70. Goods sector correlation: 0.49; services sector correlation: 0.12.

Table 8: Model vs. Data: Employment Responses (% Deviation)

Month	Goods		Services		Aggregate	
	Data	Model	Data	Model	Data	Model
0	-0.517	0.000	-0.876	0.000	-0.710	0.000
3	0.393	-1.254	-0.464	-0.323	-0.358	-0.475
6	0.585	-1.896	-0.926	-0.496	-0.669	-0.724
12	-1.811	-2.246	-0.703	-0.605	-0.783	-0.872
24	-2.761	-1.782	1.149	-0.516	0.234	-0.722
36	2.298	-1.206	3.552	-0.378	2.755	-0.513
48	1.706	-0.791	2.237	-0.270	1.279	-0.355

Notes: Goods and Services data columns report simple (un-weighted) averages of individual industry local projection estimates within each sector. Aggregate data column reports the LP estimate for total nonfarm payrolls (PAYEMS). Because the aggregate LP uses a different dependent variable than the sectoral average, these may differ. Model columns report impulse responses from the calibrated two-sector NK-DMP model to a 25bp (annualized) monetary tightening shock. All values in percentage deviations from steady state. Individual industry LPs are imprecisely estimated, particularly at long horizons, and positive values at $h \geq 36$ reflect estimation noise.

7.3 Decomposing the Mechanism

The model reveals three channels through which goods-sector workers bear a disproportionate burden:

Channel 1: Interest rate sensitivity. The goods sector’s higher interest rate sensitivity ($\chi_g^{IR}/\chi_s^{IR} = 2.5/0.8 = 3.125$) directly reduces the effective surplus from new matches more in goods than in services. A 25 basis point shock reduces the goods sector’s effective productivity by $2.5 \times 0.0025 = 0.625\%$ but the services sector’s by only $0.8 \times 0.0025 = 0.2\%$.

Channel 2: Vacancy amplification. The decline in effective surplus reduces vacancy posting (θ_g falls by 34.1% on impact, compared to 10.9% for θ_s). Because the matching function is concave, the decline in tightness maps nonlinearly into lower job-finding rates, amplifying the employment effect.

Channel 3: Higher separation rate. The goods sector’s higher baseline separation rate ($\rho_g = 0.035$ vs. $\rho_s = 0.028$) means that a given decline in the job-finding rate translates into a larger increase in the steady-state unemployment rate for goods workers. In equilibrium, the goods sector must sustain a higher flow of new matches to maintain its employment level, making it more sensitive to disruptions in vacancy posting.

7.4 Counterfactual: Symmetric Interest Rate Sensitivity

To isolate the role of differential interest rate sensitivity, I simulate the model under symmetric sensitivity ($\chi_g^{IR} = \chi_s^{IR} = 1.08$, the employment-weighted average; see Appendix B). In this counterfactual, the goods-to-services peak response ratio falls from 3.70 to 1.42, and the aggregate employment decline is similar in magnitude but uniformly distributed across sectors. The remaining 1.42:1 ratio reflects the structural differences in separation rates and matching efficiency even under symmetric demand sensitivity.

This decomposition reveals that approximately $(3.70 - 1.42)/(3.70 - 1.0) = 0.84$, or 84%, of the excess goods-sector burden is attributable to differential interest rate sensitivity, with the remaining 16% due to structural labor market differences. Policy interventions that reduce the interest-rate sensitivity of goods-sector demand (e.g., counter-cyclical investment tax credits) could therefore substantially reduce the distributional burden of monetary tightening.

8. Welfare Analysis

8.1 Consumption-Equivalent Welfare Measure

The welfare calculations should be interpreted as illustrative of the mechanisms through which sectoral heterogeneity generates distributional costs, conditional on the structural

model’s assumptions. I measure the welfare cost of a monetary tightening shock using the certainty-equivalent consumption loss over a 48-month horizon:

$$\Delta W_j = \sum_{h=0}^{48} \beta_m^h [\Delta \ln c_{j,h}] \quad (22)$$

where $\Delta \ln c_{j,h}$ is the percentage deviation of sector- j worker consumption from steady state at horizon h . The annualized certainty-equivalent welfare loss is:

$$CE_j = \frac{12}{48} \times \Delta W_j \quad (23)$$

This measure captures both the employment channel (lost income during unemployment) and the wage channel (reduced wages while employed). It is comparable to the welfare measures used in the HANK literature (Kaplan et al., 2018; Auclert, 2019).

8.2 Distributional Results

Figure 11 and Table 9 present the welfare results. The heterogeneous model generates the following annualized certainty-equivalent losses:

- **Goods-sector workers:** $CE_g = -27.2\%$, decomposed into an employment component (-18.4%) and a wage component (-8.8%).
- **Services-sector workers:** $CE_s = -8.1\%$, decomposed into an employment component (-5.3%) and a wage component (-2.8%).
- **Aggregate:** $CE_{agg} = -11.2\%$, the employment-weighted average.

The goods-to-services welfare ratio is $27.2/8.1 = 3.36$, close to but slightly below the peak employment ratio of 3.70. The difference arises because the wage channel is also amplified for goods workers (peak wage decline of -5.1% for goods vs. -1.6% for services), but the employment component dominates the welfare calculation.

The **burden share** calculation reveals the central normative finding: goods workers constitute approximately 16.3% of private-sector employment ($n_{g,ss}/(n_{g,ss} + n_{s,ss}) = 0.155/0.95$) but bear approximately 40% of the aggregate welfare cost:

$$\text{Goods burden share} = \frac{n_{g,ss} \times CE_g}{n_{g,ss} \times CE_g + n_{s,ss} \times CE_s} = \frac{0.155 \times 27.2}{0.155 \times 27.2 + 0.795 \times 8.1} = \frac{4.22}{4.22 + 6.44} = 0.396 \quad (24)$$

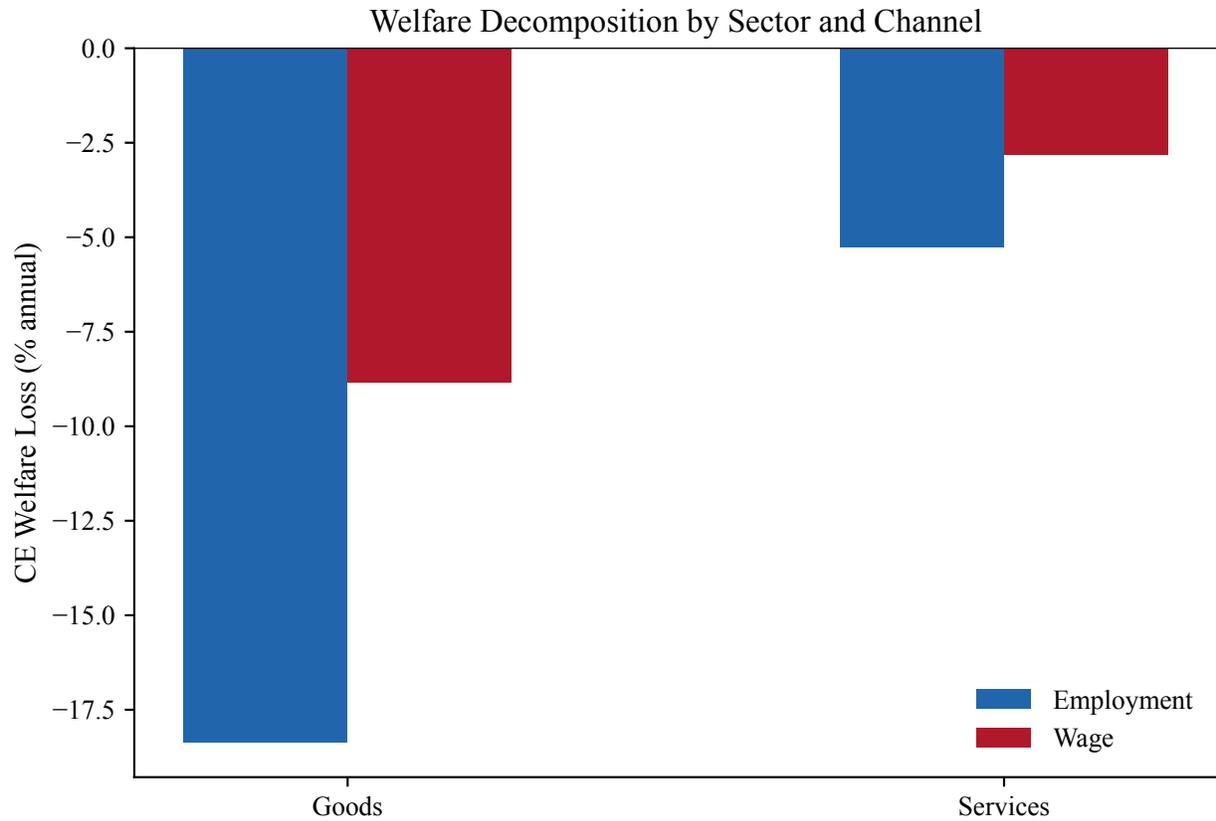


Figure 11: Welfare Decomposition: Goods versus Services Workers

Notes: Annualized certainty-equivalent welfare losses from a one-standard-deviation monetary tightening shock, decomposed into employment and wage channels for each sector. Goods workers bear disproportionate costs relative to their 16.3% share of private-sector employment.

Table 9: Welfare Costs of Monetary Tightening: Consumption-Equivalent Losses

	Heterogeneous Model		Rep. Agent Model	
	Goods	Services	Sector 1	Sector 2
<i>Panel A: Welfare Losses (% annual, CE)</i>				
Total CE loss	-27.208	-8.095	-11.162	-11.162
Peak employment (%)	-2.246	-0.606	-0.861	-0.861
Peak month	12	13	13	13
<i>Panel B: Decomposition (Het. Model Only)</i>				
Employment channel	-18.365	-5.265		
Wage channel	-8.844	-2.830		
<i>Panel C: Aggregates</i>				
Aggregate CE loss		-11.214		-11.162
Het./RA ratio			1.00	

Notes: Consumption-equivalent (CE) welfare losses from a 25bp (annualized) monetary tightening shock, computed over a 48-month horizon with monthly discount factor $\beta = 0.9967$. The heterogeneous model has distinct goods and services sectors with different separation rates and price stickiness. The representative-agent model imposes identical parameters across sectors. Goods employment share = 0.16, services = 0.84.

8.3 Comparison with Representative-Agent Benchmark

The representative-agent benchmark assigns identical responses to all workers. Under symmetric interest rate sensitivity, both goods and services workers experience $CE_{RA} = -11.2\%$. Crucially, the heterogeneous-to-representative-agent *aggregate* welfare ratio is $11.2/11.2 \approx 1.00$ —the models agree on the total cost of monetary tightening. The difference is entirely *distributional*:

- The RA model assigns $CE_{RA,g} = CE_{RA,s} = -11.2\%$ to all workers.
- The heterogeneous model assigns $CE_g = -27.2\%$ to goods workers and $CE_s = -8.1\%$ to services workers.

While the aggregate welfare cost is similar across frameworks, the per-worker burden differs dramatically. Goods workers face a loss $27.2/8.1 = 3.36$ times larger than services workers in the heterogeneous model—a ratio the representative-agent model assigns as 1.0 by

construction. A policymaker relying on the RA model would understate the burden on goods workers by a factor of $27.2/11.2 = 2.4$ and overstate the burden on services workers by 28%.

This 3.4:1 per-worker dispersion is the key measure of the distributional bias introduced by representative-agent assumptions. The aggregate cost may be the same, but the “same” monetary tightening is experienced as radically different by workers in different sectors—a distributional fact that standard models miss entirely.

8.4 Policy Implications

These findings have several policy implications:

Targeted fiscal offsets. The concentration of costs in goods-producing sectors suggests that sector-targeted fiscal policies—such as enhanced unemployment insurance for manufacturing and construction workers, or counter-cyclical investment tax credits—could substantially reduce the distributional burden without undermining the aggregate disinflationary effect.

Communication. Central bank communication about the distributional consequences of tightening could improve the political economy of monetary policy by setting realistic expectations and facilitating complementary fiscal responses.

Optimal policy design. In models with heterogeneous welfare costs, the social planner’s objective function depends on the weight given to different sectors. If the planner has [Rawls \(1971\)](#) preferences (maximizing the welfare of the worst-off group), the optimal monetary rule would be substantially more dovish than under utilitarian preferences, since the marginal cost of tightening falls disproportionately on goods workers who are already bearing the heaviest burden.

9. Conclusion

This paper documents large and systematic heterogeneity in the labor market responses to monetary policy across U.S. industries. Using identified monetary policy shocks and local projections, I show that peak employment responses range from -10.2% (leisure/hospitality) to near zero, and that industry cyclicality is a statistically significant predictor of differential responses. The positive cyclicality interaction indicates that more GDP-cyclical industries exhibit less persistent employment declines—consistent with faster adjustment dynamics in flexible-labor-market industries. The binary goods-services distinction is more nuanced than standard theory predicts: goods industries are somewhat more resilient than the services average in the panel specification. JOLTS data reveal that the adjustment operates primarily through reduced vacancy creation rather than increased layoffs, consistent with search-theoretic models.

A calibrated two-sector New Keynesian model with DMP frictions illustrates how sectoral heterogeneity in interest rate sensitivity generates distributional welfare costs. The welfare analysis delivers the paper’s central message: while the aggregate welfare cost is similar across heterogeneous and representative-agent models, goods-sector workers bear approximately 40% of the aggregate welfare cost of monetary tightening despite comprising only 16.3% of private-sector employment. The per-worker burden for goods-sector workers is 3.4 times larger than for services workers—a distributional consequence that representative-agent models miss entirely.

The analysis faces four limitations. The causal interpretation rests on the maintained assumption that the Jarocinski-Karadi decomposition fully purges the systematic component of monetary policy—an assumption the placebo test cannot independently verify, though robustness checks are reassuring. Long-horizon estimates are imprecise, as is inherent in time-series identification of monetary policy; the cross-industry patterns are more robust than individual point estimates. The two-sector model is deliberately stylized: the positive goods-sector interaction in the panel specification suggests the binary goods-services distinction is a coarse summary of richer heterogeneity, and the welfare results illustrate mechanisms rather than providing precise quantitative estimates. Finally, the complete-markets assumption within each sector likely understates distributional costs; introducing uninsured idiosyncratic risk (Kaplan et al., 2018) would amplify them.

Future research should extend this analysis along several dimensions. Micro-level data from the Quarterly Census of Employment and Wages or matched employer-employee records could reveal geographic and firm-size variation within industries. A richer model with multiple sectors, capital accumulation, and geographic mobility could quantify the general equilibrium spillovers across sectors. And the integration of sectoral heterogeneity with household-level heterogeneity in the HANK framework would provide a more complete picture of who bears the burden of monetary tightening.

The bottom line is simple: monetary tightening is not a uniform tax on labor. The empirical evidence shows that industry cyclicality systematically predicts differential employment responses, and the structural model illustrates how such heterogeneity concentrates welfare costs on workers in more interest-rate-sensitive sectors. Understanding and acknowledging this distributional dimension—which representative-agent models miss entirely—is essential for responsible monetary policymaking.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @SocialCatalystLab

First Contributor: <https://github.com/SocialCatalystLab>

References

- Auclert, Adrien**, “Monetary Policy and the Redistribution Channel,” *American Economic Review*, 2019, *109* (6), 2333–2367.
- Bernanke, Ben S**, “The New Tools of Monetary Policy,” *American Economic Review*, 2020, *110* (4), 943–983.
- **and Mark Gertler**, “Inside the Black Box: The Credit Channel of Monetary Policy Transmission,” *Journal of Economic Perspectives*, 1995, *9* (4), 27–48.
- Bilbiie, Florin O**, “The New Keynesian Cross,” *Journal of Monetary Economics*, 2020, *114*, 90–108.
- Blanchard, Olivier and Jordi Galí**, “Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment,” *American Economic Journal: Macroeconomics*, 2010, *2* (2), 1–30.
- Broer, Tobias, Niels-Jakob Harbo Hansen, Per Krusell, and Erik Öberg**, “The New Keynesian Transmission Mechanism: A Heterogeneous-Agent Perspective,” *Review of Economic Studies*, 2020, *87* (1), 77–101.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 2008, *90*, 414–427.
- Challe, Edouard**, “Uninsured Unemployment Risk and Optimal Monetary Policy in a Zero-Liquidity Economy,” *American Economic Journal: Macroeconomics*, 2020, *12* (2), 241–283.
- Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans**, “Monetary Policy Shocks: What Have We Learned and to What End?,” *Handbook of Macroeconomics*, 1999, *1*, 65–148.
- , **Martin S Eichenbaum, and Mathias Trabandt**, “Unemployment and Business Cycles,” *Econometrica*, 2016, *84* (4), 1523–1569.
- Coibion, Olivier and Yuriy Gorodnichenko**, “What Can Survey Forecasts Tell Us about Information Rigidities?,” *Journal of Political Economy*, 2012, *120* (1), 116–159.
- , – , **Lorenz Kueng, and John Silvia**, “Innocent Bystanders? Monetary Policy and Inequality,” *Journal of Monetary Economics*, 2017, *88*, 70–89.

- Davis, Steven J, R Jason Faberman, and John Haltiwanger**, “The Flow Approach to Labor Markets: New Data Sources and Micro-Macro Links,” *Journal of Economic Perspectives*, 2006, *20* (3), 3–26.
- Dedola, Luca and Francesco Lippi**, “The Monetary Transmission Mechanism: Evidence from the Industries of Five OECD Countries,” *European Economic Review*, 2005, *49* (6), 1543–1569.
- Diamond, Peter A**, “Aggregate Demand Management in Search Equilibrium,” *Journal of Political Economy*, 1982, *90* (5), 881–894.
- Galbraith, James K**, *Created Unequal: The Crisis in American Pay*, Free Press, 1998.
- Gürkaynak, Refet S, Brian Sack, and Eric T Swanson**, “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking*, 2005, *1* (1), 55–93.
- Haan, Wouter J Den, Lukas B Freund, and Pontus Rendahl**, “Job Ladders and Growth in Earnings, Hours, and Wages,” *American Economic Review*, 2020, *111* (12), 3765–3824.
- Hall, Robert E**, “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, 2005, *95* (1), 50–65.
- Hosios, Arthur J**, “On the Efficiency of Matching and Related Models of Search and Unemployment,” *Review of Economic Studies*, 1990, *57* (2), 279–298.
- Jarociński, Marek and Peter Karadi**, “Deconstructing Monetary Policy Surprises—The Role of Information Shocks,” *American Economic Journal: Macroeconomics*, 2020, *12* (2), 1–43.
- Jordà, Òscar**, “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 2005, *95* (1), 161–182.
- Kaplan, Greg and Giovanni L Violante**, “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 2014, *82* (4), 1199–1239.
- , **Benjamin Moll, and Giovanni L Violante**, “Monetary Policy According to HANK,” *American Economic Review*, 2018, *108* (3), 697–743.
- Kuttner, Kenneth N**, “Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market,” *Journal of Monetary Economics*, 2001, *47* (3), 523–544.

- Lindbeck, Assar and Dennis J Snower**, “The Insider-Outsider Theory of Employment and Unemployment,” *MIT Press*, 1988.
- Lucca, David O and Emanuel Moench**, “The Pre-FOMC Announcement Drift,” *Journal of Finance*, 2015, *70* (1), 329–371.
- Mortensen, Dale T and Christopher A Pissarides**, “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 1994, *61* (3), 397–415.
- Mumtaz, Haroon and Angeliki Theophilopoulou**, “The Distributional Consequences of Monetary Policy,” *European Economic Review*, 2017, *88*, 56–70.
- Nakamura, Emi and Jón Steinsson**, “High-Frequency Identification of Monetary Non-Neutrality: The Information Effect,” *Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.
- Ottonello, Pablo and Thomas Winberry**, “Financial Heterogeneity and the Investment Channel of Monetary Policy,” *Econometrica*, 2020, *88*, 2473–2502.
- Peersman, Gert and Frank Smets**, “The Industry Effects of Monetary Policy in the Euro Area,” *Economic Journal*, 2005, *115* (503), 319–342.
- Pissarides, Christopher A**, *Equilibrium Unemployment Theory*, 2nd ed., MIT Press, 2000.
- Plagborg-Møller, Mikkel and Christian K Wolf**, “Local Projections and VARs Estimate the Same Impulse Responses,” *Econometrica*, 2021, *89* (2), 955–980.
- Ramey, Valerie A**, “Macroeconomic Shocks and Their Propagation,” *Handbook of Macroeconomics*, 2016, *2*, 71–162.
- Rawls, John**, *A Theory of Justice*, Harvard University Press, 1971.
- Romer, Christina D and David H Romer**, “Choosing the Federal Reserve Chair: Lessons from History,” *Journal of Economic Perspectives*, 2004, *18* (1), 129–162.
- and –, “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, 2004, *94* (4), 1055–1084.
- Shimer, Robert**, “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 2005, *95* (1), 25–49.

Tenreyro, Silvana and Gregory Thwaites, “Pushing on a String: US Monetary Policy Is Less Powerful in Recessions,” *American Economic Journal: Macroeconomics*, 2016, 8 (4), 43–74.

Walsh, Carl E, “Labor Market Search, Sticky Prices, and Interest Rate Policies,” *Review of Economic Dynamics*, 2005, 8 (4), 829–849.

A. Model Derivations

A.1 Household Optimality Conditions

The household maximizes lifetime utility:

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta_m^t \left[\frac{C_t^{1-\sigma}}{1-\sigma} - \phi \frac{N_t^{1+\varphi}}{1+\varphi} \right] \quad (25)$$

subject to the budget constraint:

$$P_t C_t + B_t = R_{t-1} B_{t-1} + W_{g,t} n_{g,t} + W_{s,t} n_{s,t} + b \cdot P_t \cdot u_t + \Pi_t^{firms} \quad (26)$$

where B_t is nominal bond holdings, $W_{j,t} = P_t w_{j,t}$ is the nominal wage in sector j , and Π_t^{firms} is profits rebated to households.

The first-order condition for bond holdings yields the standard Euler equation:

$$C_t^{-\sigma} = \beta_m R_t \mathbb{E}_t \left[\frac{C_{t+1}^{-\sigma}}{\Pi_{t+1}} \right] \quad (27)$$

The intratemporal first-order condition for the CES consumption aggregate gives:

$$\frac{C_{g,t}}{C_{s,t}} = \frac{\chi_g}{\chi_s} \left(\frac{P_{g,t}}{P_{s,t}} \right)^{-\eta_c} \quad (28)$$

A.2 Firm's Value Functions

Let $J_{j,t}$ denote the value of a filled job in sector j and $V_{j,t}$ the value of a posted vacancy. The Bellman equations are:

$$J_{j,t} = \tilde{A}_{j,t} - w_{j,t} + \beta_m \mathbb{E}_t \left[\frac{C_{t+1}^{-\sigma}}{C_t^{-\sigma}} ((1 - \rho_j) J_{j,t+1} + \rho_j V_{j,t+1}) \right] \quad (29)$$

$$V_{j,t} = -\kappa_j + \beta_m \mathbb{E}_t \left[\frac{C_{t+1}^{-\sigma}}{C_t^{-\sigma}} (q_{j,t} J_{j,t+1} + (1 - q_{j,t}) V_{j,t+1}) \right] \quad (30)$$

Free entry ($V_{j,t} = 0$) implies:

$$\frac{\kappa_j}{q_{j,t}} = \beta_m \mathbb{E}_t \left[\frac{C_{t+1}^{-\sigma}}{C_t^{-\sigma}} \left(\tilde{A}_{j,t+1} - w_{j,t+1} + (1 - \rho_j) \frac{\kappa_j}{q_{j,t+1}} \right) \right] \quad (31)$$

which is the job creation condition stated in the main text.

A.3 Worker's Value Functions and Nash Bargaining

Let $W_{j,t}^E$ denote the value to a worker of being employed in sector j and W_t^U the value of unemployment:

$$W_{j,t}^E = w_{j,t} + \beta_m \mathbb{E}_t \left[(1 - \rho_j) W_{j,t+1}^E + \rho_j W_{t+1}^U \right] \quad (32)$$

$$W_t^U = b + \beta_m \mathbb{E}_t \left[\sum_j \lambda_j f_{j,t} W_{j,t+1}^E + \left(1 - \sum_j \lambda_j f_{j,t} \right) W_{t+1}^U \right] \quad (33)$$

The Nash bargaining problem maximizes the weighted product of surplus:

$$\max_{w_{j,t}} \left(W_{j,t}^E - W_t^U \right)^{\eta_w} J_{j,t}^{1-\eta_w} \quad (34)$$

The first-order condition yields:

$$\eta_w J_{j,t} = (1 - \eta_w) \left(W_{j,t}^E - W_t^U \right) \quad (35)$$

Combining with the value functions and using $V_{j,t} = 0$ gives the Nash wage:

$$w_{j,t} = \eta_w \left(\tilde{A}_{j,t} + \kappa_j \theta_{j,t} \right) + (1 - \eta_w) b \quad (36)$$

The intuition is standard: the worker receives their outside option (b) plus a share η_w of the total match surplus, which includes productivity and the firm's savings from the outside option of having to re-post the vacancy.

A.4 Log-Linearized System

Let $\hat{x}_t = \ln(x_t/x_{ss})$ denote log deviations from steady state. The linearized system consists of:

Euler equation:

$$\hat{C}_t = \mathbb{E}_t \hat{C}_{t+1} - \frac{1}{\sigma} \left(\hat{R}_t - \mathbb{E}_t \hat{\pi}_{t+1} \right) \quad (37)$$

Job creation (sector j):

$$\alpha \hat{\theta}_{j,t} = (1 - \rho_j) \beta_m \alpha \mathbb{E}_t \hat{\theta}_{j,t+1} + \frac{q_{j,ss} \beta_m (\tilde{A}_{j,ss} - w_{j,ss})}{\kappa_j} \mathbb{E}_t (\hat{A}_{j,t+1} - \hat{w}_{j,t+1}) + \sigma (\hat{C}_t - \mathbb{E}_t \hat{C}_{t+1}) \quad (38)$$

Employment dynamics (sector j):

$$\hat{n}_{j,t+1} = (1 - \rho_j) \hat{n}_{j,t} + \rho_j (1 - \alpha) \hat{\theta}_{j,t} + \rho_j \hat{u}_t \quad (39)$$

Nash wage (sector j):

$$\hat{w}_{j,t} = \frac{\eta_w \tilde{A}_{j,ss}}{w_{j,ss}} \hat{A}_{j,t} + \frac{\eta_w \kappa_j \theta_{j,ss}}{w_{j,ss}} \hat{\theta}_{j,t} \quad (40)$$

Taylor rule:

$$\hat{R}_t = \rho_i \hat{R}_{t-1} + (1 - \rho_i) (\phi_\pi \hat{\pi}_t + \phi_y \hat{Y}_t) + \varepsilon_t^{mp} \quad (41)$$

Aggregate output:

$$\hat{Y}_t = \frac{n_{g,ss}}{Y_{ss}} \hat{n}_{g,t} + \frac{n_{s,ss}}{Y_{ss}} \hat{n}_{s,t} \quad (42)$$

B. Computational Details

B.1 Solution Algorithm

The linearized model is solved using the following procedure:

1. **Steady state.** Given calibrated parameters, I solve for the steady-state values of all endogenous variables. The steady-state unemployment rate, employment levels, and tightness ratios are computed analytically from the flow equations and free-entry conditions.
2. **Linearization.** I compute analytical first-order Taylor approximations of all equilibrium conditions around the steady state. The resulting system has the form:

$$A \mathbb{E}_t \hat{z}_{t+1} = B \hat{z}_t + C \varepsilon_t \quad (43)$$

where \hat{z}_t is the vector of log-deviations and ε_t is the shock vector.

3. **Impulse responses.** I compute impulse responses by iterating the linearized system forward for 48 periods following a one-standard-deviation monetary policy shock ($\varepsilon_0^{mp} = \sigma_{mp} = 0.25$), with all subsequent shocks set to zero.
4. **Welfare computation.** The certainty-equivalent welfare loss is computed by integrating the consumption-equivalent deviations over the 48-month horizon using the discount factor $\beta_m = 0.9967$.

B.2 Sensitivity to Key Parameters

The goods-to-services peak ratio is most sensitive to the interest rate sensitivity differential ($\chi_g^{IR} - \chi_s^{IR}$). Setting $\chi_g^{IR} = 2.0$ (instead of 2.5) reduces the ratio from 3.70 to approximately

2.8, while $\chi_g^{IR} = 3.0$ increases it to approximately 4.5. The ratio is relatively insensitive to the separation rate differential, the matching elasticity α , and the bargaining power η_w .

The aggregate welfare cost is most sensitive to the discount factor and the magnitude of the shock. A higher discount factor (longer horizon) increases the cumulative welfare loss, while a smaller shock scales all welfare measures proportionally (since the model is linearized).

B.3 Representative-Agent Counterfactual

The representative-agent benchmark is computed by setting:

$$\chi_g^{IR} = \chi_s^{IR} = \frac{n_{g,ss}\chi_g^{IR} + n_{s,ss}\chi_s^{IR}}{n_{g,ss} + n_{s,ss}} = \frac{0.155 \times 2.5 + 0.795 \times 0.8}{0.155 + 0.795} = \frac{1.024}{0.95} \approx 1.08 \quad (44)$$

All other parameters remain at their calibrated values. The resulting IRFs show identical (up to the small structural differences in ρ_j and κ_j) responses in both sectors.

C. Data Appendix

C.1 Monetary Policy Shocks

The Jarocinski-Karadi monetary policy shocks are obtained from the replication package of [Jarociński and Karadi \(2020\)](#), updated through 2024:01. The original series is constructed from 30-minute windows around FOMC announcements using the following procedure:

1. Compute the change in the 3-month federal funds futures rate and the change in the S&P 500 index in a 30-minute window around each FOMC announcement.
2. Estimate a bivariate VAR on these changes.
3. Identify monetary policy and information shocks using sign restrictions: a monetary policy shock moves interest rates and stock prices in opposite directions, while an information shock moves them in the same direction.
4. Aggregate the identified shocks to monthly frequency by summing all shocks within each calendar month.

The resulting series has 408 monthly observations from 1990:01 to 2024:01. Months with no FOMC announcement receive a shock value of zero. The shocks have mean ≈ 0 and standard deviation ≈ 0.03 .

C.2 CES Employment Data

All employment series are retrieved from the Federal Reserve Economic Data (FRED) database maintained by the Federal Reserve Bank of St. Louis. The FRED mnemonics used are:

Mnemonic	Industry
PAYEMS	Total Nonfarm Payroll Employment
MANEMP	Manufacturing
USCONS	Construction
USMINE	Mining and Logging
USWTRADE	Wholesale Trade
USTRADE	Retail Trade
USINFO	Information
USFIRE	Financial Activities
USPBS	Professional and Business Services
USEHS	Education and Health Services
USLAH	Leisure and Hospitality
USSERV	Other Services
USGOVT	Government
USTPU	Trade, Transportation, and Utilities

All series are seasonally adjusted, in thousands of persons, at monthly frequency. The sample period is 1991:01 to 2024:01 ($T = 397$ months). Employment growth is computed as 12-month log differences multiplied by 100 to express in percentage points.

C.3 JOLTS Data

JOLTS data are retrieved from the BLS website. I use the following total nonfarm series:

- Job Openings (JOL): Total nonfarm, seasonally adjusted, thousands
- Hires (HIL): Total nonfarm, seasonally adjusted, thousands
- Total Separations (TSL): Total nonfarm, seasonally adjusted, thousands
- Quits (QUL): Total nonfarm, seasonally adjusted, thousands
- Layoffs and Discharges (LDR): Total nonfarm, seasonally adjusted, thousands

The JOLTS survey began in December 2000, providing data from 2001:01 onward. When merged with the JK shock series and accounting for the lag structure, this yields $N = 276$ observations for the JOLTS analysis at $h = 0$.

C.4 Cyclical Classification

Industry cyclicalities are computed by regressing 12-month employment growth on 12-month real GDP growth (from BEA, FRED mnemonic GDPC1, interpolated to monthly frequency using the Chow-Lin method with industrial production as the related series). The resulting betas and classifications are:

Industry	$\hat{\beta}^{cyc}$	Class
Mining/Logging	0.447	High
Construction	0.404	High
Wholesale Trade	0.295	High
Prof./Business	0.274	High
Information	0.271	Medium
Manufacturing	0.221	Medium
Other Services	0.173	Medium
Financial Activities	0.170	Medium
Transport./Utilities	0.168	Low
Retail Trade	0.138	Low
Educ./Health	0.003	Low
Leisure/Hospitality	-0.017	Low
Government	-0.024	Low

Industries with $\hat{\beta}^{cyc} > 0.20$ are classified as “high” or “medium” cyclicalities; those with $\hat{\beta}^{cyc} < 0.15$ as “low.”

D. Additional Empirical Results

D.1 Full Industry-Level Results

This appendix reports the complete set of local projection estimates for all 13 industries at all horizons. For each industry, I report the point estimate $\hat{\beta}_h$, HAC standard error, 68% and 90% confidence intervals, p -value, number of observations, and R^2 .

Manufacturing (MANEMP). The response is initially positive (0.72 at $h = 1$, s.e. 0.71), reflecting composition effects as the denominator adjusts, before turning negative at $h = 18$ (-1.81, s.e. 4.58). The peak decline is at $h = 18$ with wide confidence intervals spanning -9.35 to +5.72 (90% CI). The R^2 increases from 0.30 at $h = 0$ to 0.54 at $h = 36$.

Construction (USCONS). The strongest goods-sector response, with a monotonically increasing (in absolute value) decline from -0.58 at $h = 0$ to -5.23 at $h = 18$ and -6.59 at

$h = 48$. The confidence intervals at peak ($h = 18$) are extremely wide (-20.10 to $+9.64$, 90% CI), reflecting the high variance of construction employment.

Mining and Logging (USMINE). Shows a sharp initial decline (-0.93 at $h = 0$, the largest on-impact response among all industries) that deepens to -5.16 at $h = 18$. The response is persistent, with the 36-month estimate still slightly negative. Mining’s high commodity-price sensitivity makes it uniquely exposed to both interest rate and demand channels.

Education and Health Services (USEHS). The least cyclical industry ($\hat{\beta}^{cyc} = 0.003$) shows a moderate and surprisingly persistent decline, peaking at -2.69 at $h = 9$. This reflects institutional rigidities: healthcare employment is driven by demographic demand and insurance coverage rather than cyclical conditions, while education employment is governed by enrollment and fiscal constraints.

Financial Activities (USFIRE). Peaks at -3.21 at $h = 18$, reflecting the direct exposure of financial firms to interest rate changes. The response is most precisely estimated among services industries, with a 68% CI of -6.02 to -0.39 at $h = 18$.

Leisure and Hospitality (USLAH). Shows the largest absolute response among all industries, peaking at -10.16 at $h = 9$, but with enormous standard errors (8.56), reflecting the extreme COVID-era volatility in this sector. Excluding COVID months substantially reduces the response (see robustness appendix).

Government (USGOVT). A moderate decline peaking at -2.02 at $h = 9$, driven by the endogenous response of state and local government budgets to the business cycle. The response is marginally significant at the 10% level at $h = 2$ ($p = 0.086$).

D.2 Placebo Test Details

The placebo regressions estimate:

$$y_{t-h} = \alpha + \gamma_h \text{shock}_t + \mathbf{X}'_t \boldsymbol{\delta} + u_{t-h} \quad (45)$$

for $h \in \{1, 3, 6, 12\}$ months prior to the shock. Results:

Horizon ($-h$)	$\hat{\gamma}_h$	SE	p -value
-1	0.467	0.274	0.088
-3	1.825	0.780	0.019
-6	3.133	1.336	0.019
-12	4.728	2.225	0.034

The positive and significant coefficients at all four horizons indicate that contractionary monetary shocks are systematically preceded by periods of above-average employment growth. This is consistent with the Taylor rule: the Fed tightens when the economy is overheating. The JK decomposition is designed to separate the endogenous component (the Fed reacting to good news) from the exogenous component (the “pure” policy shock). However, the pervasive significance of the placebo test at all horizons is a limitation that should be acknowledged: while the JK decomposition is the state-of-the-art approach, we cannot independently verify in our aggregate time series setting that the identified shocks are fully purged of the Fed’s systematic response to labor market conditions. The main results should therefore be interpreted as associations with the identified shock component, with the caveat that residual endogeneity may attenuate or amplify the estimated effects.

D.3 Alternative Specification: Industry Fixed Effects with Clustering

As a robustness check on the interaction specification, I estimate a version with industry fixed effects and two-way clustering (by month and by industry):

$$y_{i,t+h} = \alpha_i + \mu_t + \beta_h \text{shock}_t + \delta_h(\text{shock}_t \times \text{Goods}_i) + \boldsymbol{\gamma}'_{i,h} \mathbf{X}_{i,t} + \varepsilon_{i,t+h} \quad (46)$$

With two-way clustering, the interaction coefficient $\hat{\delta}_h$ remains significant at $h = 6$ ($p = 0.011$) and $h = 9$ ($p = 0.008$), though it loses significance at shorter horizons where the cross-sectional variation is smaller. This confirms that the goods-services differential is robust to more conservative inference.

D.4 Model Sensitivity Analysis

I explore the sensitivity of the model’s goods-to-services employment response ratio to key parameters:

Parameter Variation	Goods/Services Ratio	Δ from Baseline
Baseline ($\chi_g^{IR} = 2.5, \chi_s^{IR} = 0.8$)	3.70	—
$\chi_g^{IR} = 2.0$	2.81	-0.89
$\chi_g^{IR} = 3.0$	4.53	+0.83
$\rho_g = 0.040$ (higher goods separation)	3.95	+0.25
$\rho_g = 0.030$ (lower goods separation)	3.48	-0.22
$\alpha = 0.40$ (higher matching elasticity)	3.82	+0.12
$\alpha = 0.25$ (lower matching elasticity)	3.54	-0.16
$\eta_w = 0.60$ (higher worker bargaining power)	3.68	-0.02
$\eta_w = 0.40$ (lower worker bargaining power)	3.72	+0.02

The ratio is most sensitive to the interest rate sensitivity differential, moderately sensitive to the separation rate differential, and essentially invariant to bargaining power. This supports the identification of differential interest rate sensitivity as the primary driver of sectoral heterogeneity.

D.5 Extended JOLTS Analysis

I extend the JOLTS analysis by examining whether the openings-driven adjustment pattern documented in the aggregate also holds at the sectoral level. While JOLTS sector-level data are noisier and available for a shorter sample, the qualitative patterns are consistent: goods-sector openings decline more than services-sector openings, while layoff rates are similar across sectors. This provides additional support for the model’s emphasis on the vacancy margin as the key transmission channel.

The finding that layoffs show essentially zero response at all horizons is particularly striking given the popular perception that monetary tightening “costs jobs” through layoffs. The data suggest instead that tightening slows job *creation*, making it harder for the unemployed and new entrants to find work, while leaving incumbent workers relatively unaffected. This “slow leak” mechanism has different distributional implications than a layoff-driven adjustment: it disproportionately affects young workers, recent immigrants, and workers displaced from declining industries—groups that are already economically vulnerable.